

Quantum Technologies

Ninth edition - 2026

Part 1/5 – Prologue

Olivier Ezratty



the lab

le lab quantique

Le Lab Quantique is a supporter and promoter of this book, but not its publisher or editor.



Understanding Quantum Technologies

Ninth edition - 2026 - Version 9.0

Part 1/5 – Prologue

Olivier Ezratty

Generated on September 28, 2026

About the author

Olivier Ezratty



0000-0003-3944-2896

olivier (at) oezratty.net
www.oezratty.net, @olivez
+33 6 67 37 92 41

Olivier Ezratty is an author, academic, teacher and advisor on quantum technologies. He started this endeavour in 2018. He is the author of the reference book “Understanding Quantum Technologies” (September 2021, 2022, 2023, 2024, 2025 and 2026) following three previous editions in French in 2018, 2019 and 2020. The 2025 edition is also available in paperback version on Amazon. This 2026 edition’s Amazon print version will be available before the end of 2026.

Since 2021, he has been teaching quantum technologies at EPITA, an IT engineering school in France, at Ecole Normale Supérieure Paris-Saclay since 2024 and CentraleSupélec since 2005. He’s affiliated with the EPITA Research lab for his various research writings.

He is a cofounder of the Quantum Energy Initiative in 2022 with Alexia Auffèves (CNRS MajuLab Singapore), Robert Whitney (CNRS LPMMC) and Janine Splettstoesser (Chalmers University, Sweden), and member of its scientific board.

He acts as an independent expert for Bpifrance and at the Agence Nationale de Recherche (ANR) to evaluate quantum R&D projects, as a member of the European Quantum Computing Advisory Board (Q-CAB) since April 2026, and of the Inria expert group for the selection of risky research projects (Programme PIQ) since June 2024.

Since 2018, he has spoken in a large number of quantum technology events such as the Q2B Paris and Santa Clara organized by QC Ware, France Quantum, Lindau Nobel Laureate meeting, DPG, Q-Expo and other events, on top of presentations at Société Générale, BNPParibas, Crédit Agricole, Michelin, Adéo, L’Oréal, FIECC, IHEDN, Business France, CentraleSupélec, Avolta Partners, IHEDN, etc.

He coproduces two series of podcasts on quantum technologies along with Fanny Bouton (in French): a monthly “Quantum ” on tech news (since September 2019) and Decode Quantum (2020-2026), with entrepreneurs and researchers, with a total of about 170 episodes as of September 2026.

He also lectures or lectured in various universities such as CentraleSupélec, ENSTA, Telecom Paristech, Les Gobelins, HEC, Neoma Rouen and SciencePo, on artificial intelligence, entrepreneurship and product management (until 2020) and on quantum technologies (since 2018), in French and English as needed. He is also the author of many open-source ebooks in French on entrepreneurship (2006-2019), the CES of Las Vegas yearly report (2006-2020) and on artificial intelligence (2016-2021).

Beforehand, Olivier Ezratty was specialized in many other topics since 2005, like digital television, internet of things and artificial intelligence. As such, he carried out various strategic advisory missions, conferences, or training assignments in different verticals and domains such as media and telecoms, finance and insurance, industry and services and the public sector.

Olivier Ezratty started in 1985 at Sogitec, a subsidiary of the Dassault group, where he was successively Software Engineer, then Head of the Research Department in the Communication Division. He initiated developments under Windows 1.0 in the field of editorial computing as well as on SGML, the ancestor of HTML and XML. Joining Microsoft France in 1990, he gained an extensive experience in many areas of the marketing mix: products, channels, markets and communication. He launched in France the first version of Visual Basic in 1991 and Windows NT in 1993. In 1998, he became Marketing and Communication Director of Microsoft France and in 2001, of the Developer Division, which he created in France to launch the .NET platform and promote it to developers, higher education and research, as well as to startups. Olivier Ezratty is a software engineer from Centrale Paris (1985), which became CentraleSupélec in 2015.

This document is provided free of charge and is licensed under a “Creative Commons” license.

in the variant “Attribution-Noncommercial-No Derivative Works 2.0”. ISSN 2680-0527 

Cover illustration credits: personal creation associating a Bloch sphere describing a qubit and the symbol of peace (my creation, first published in 2018) above a long list of terms from the book glossary. This book contains about 1,325 illustrations. I have managed to give credit to their creators as much as possible. Most sources are credited in the captions. Only scientists’ portraits are not credited since it’s quite hard to track their original source. I have added my own credit in most of the illustrations I have created or modified. In some cases, I have redrawn some third-party illustrations to create clean vector versions or used existing third-party illustrations and added my own text comments. The originals are still credited in that case.

Table of contents

Table of contents	iii	2.6.13 International collaboration	93
Foreword	vi	2.6.14 Technology Readiness Level	94
Contents	viii	2.6.15 Bibliography and notes	95
1 Why	1	2.7 Quantum physics history and scientists key takeaways	99
1.1 A domain in search of pedagogy	2	3 Quantum physics 101	100
1.2 A new technology wave	3	3.1 Postulates	101
1.3 Reading guide	4	3.2 Quantization	103
1.4 First and second quantum revolutions applications	6	3.2.1 Black body spectrum explanation ..	107
1.5 Why quantum computing?	7	3.2.2 Electrons quantum numbers	109
1.5.1 Quantum computing promise	8	3.2.3 Nucleons and nucleus quantum properties	112
1.5.2 Moore's law	10	3.2.4 Photon quantum numbers	115
1.5.3 Transistor density evolution	11	3.2.5 Elementary particles quantum numbers	120
1.5.4 Classical computing technology developments	14	3.2.6 Spins everywhere	122
1.5.5 Unconventional computing	20	3.3 Wave-particle duality	124
1.6 Why... key takeaways	22	3.3.1 Schrödinger's wave equation	125
1.7 Bibliography and notes	23	3.3.2 Delayed choice experiment	128
2 History and scientists	25	3.3.3 Large objects wave behavior	129
2.1 Precursors	27	3.3.4 Photon's wave-particle duality	129
2.2 Founders	33	3.4 Superposition and entanglement	130
2.3 Post-war	48	3.4.1 Superposition	131
2.4 Quantum technologies physicists	52	3.4.2 Entanglement	132
2.4.1 Generalists	52	3.5 Indetermination	139
2.4.2 Cold atoms	55	3.6 Measurement	141
2.4.3 Trapped ions	57	3.6.1 Continuous properties measurement ..	141
2.4.4 Superconductivity	58	3.6.2 Discrete properties measurement	141
2.4.5 Spin qubits	60	3.6.3 Qubit measurement	141
2.4.6 NV centers	61	3.6.4 Quantum speed limits	142
2.4.7 Photonics	62	3.6.5 The quantum measurement problem ..	142
2.4.8 Quantum communications and cryptography	63	3.7 No-cloning	143
2.4.9 Other academic fields	64	3.8 Tunnel effect	144
2.5 Quantum information science and algorithms creators	65	3.9 Symmetries	144
2.5.1 Theory	65	3.9.1 Bibliography and notes	146
2.5.2 Algorithms	67	3.10 Quantum matter	150
2.5.3 Error correction	69	3.10.1 Definitions	150
2.5.4 Other domains	70	3.10.2 Superconductivity	153
2.5.5 Bibliography and notes	71	3.10.3 Superfluidity	160
2.6 Research for dummies	75	3.10.4 Bose-Einstein Condensates	161
2.6.1 Long-term view	75	3.10.5 Supersolids	162
2.6.2 Scientific papers	76	3.10.6 Polaritons	162
2.6.3 Writing a paper	80	3.10.7 Magnons	165
2.6.4 Publishing a paper	83	3.10.8 Skyrmions	166
2.6.5 Communicating on a paper	85	3.10.9 Topological matter	166
2.6.6 Predatory, scam and fake conferences ..	86	3.10.10 Time crystals	167
2.6.7 Papers analysis and classification ...	86	3.10.11 Quantum batteries	169
2.6.8 The generative AI revolution	89	3.10.12 Bibliography and notes	171
2.6.9 Academic roles	92	3.11 Extreme quantum	177
2.6.10 h-index	92	3.11.1 Quantum field theory	177
2.6.11 Poster sessions	93	3.11.2 Quantum vacuum fluctuation	177
2.6.12 Figures of merit	93	3.11.3 Casimir effect	179
		3.11.4 Theories of everything	182
		3.11.5 Bibliography and notes	185
		3.12 Quantum foundations	187
		3.12.1 Quantum physics and its missing ontology	187

3.12.2	Quantum foundations questions	188	5.1.6	Architecture overview	271
3.12.3	Quantum foundations interpretations	192	5.1.7	Processor layout	275
3.12.4	Other interpretations	198	5.1.8	Bibliography and notes	277
3.12.5	Bibliography and notes	199	5.2	Error handling	280
3.13	Quantum physics 101 key takeaways	205	5.2.1	Error types and sources	280
4	Gate-based quantum computing	206	5.2.2	Qubits fidelities	286
4.1	Linear algebra	206	5.2.3	Qubit figures of merit	289
4.1.1	Linearity	206	5.2.4	Quantum error mitigation	289
4.1.2	Hilbert spaces and orthonormal bases	207	5.2.5	Quantum error suppression and detection	293
4.1.3	Dirac Notation	207	5.2.6	Error correction basics	294
4.1.4	Eigenstuff	208	5.2.7	Error correction codes	298
4.1.5	Tensor products	209	5.2.8	Error syndrome decoding	311
4.1.6	Entanglement	209	5.2.9	Logical qubits	315
4.1.7	Matrices	210	5.2.10	Magic states	320
4.1.8	Pure and mixed states	211	5.2.11	Fault-tolerant quantum computing .	323
4.1.9	Density matrices	213	5.2.12	Intermediate FTQC approaches	331
4.1.10	Grad, curls and divs	218	5.2.13	Gate-based computing classes	332
4.1.11	Permanent and determinant	218	5.2.14	FTQC resource estimations	334
4.1.12	Fourier transforms	219	5.2.15	Bibliography and notes	341
4.1.13	Lie groups	220	5.3	Quantum memory	359
4.1.14	Nonlinearities	220	5.3.1	Classical memory	359
4.2	Qubits	221	5.3.2	Quantum algorithms requirements .	359
4.2.1	Mathematical qubit	222	5.3.3	Quantum memory types	360
4.2.2	Bloch sphere	222	5.3.4	Quantum memory physical imple- mentations	363
4.2.3	Physical qubit	224	5.3.5	Bibliography and notes	367
4.2.4	Logical qubit	226	5.4	Energetics	370
4.2.5	Qutrits and qudits	226	5.4.1	ICT lessons	370
4.3	Registers	227	5.4.2	First assessments	375
4.4	Gates	230	5.4.3	Quantum computing energetics	379
4.4.1	Common gates	230	5.4.4	Qubit type differentiation	390
4.4.2	Physical gates	235	5.4.5	Actions	392
4.4.3	Gate classes and universal gate sets	237	5.4.6	Bibliography and notes	395
4.4.4	Gate synthesis	242	5.5	Lieb-Robinson bounds	402
4.5	Inputs and outputs	242	5.5.1	Probing information propagation ...	402
4.6	Qubit lifecycle	243	5.5.2	Nonlocality vs information prop- agation	404
4.7	Measurement	245	5.5.3	Vector state vs quantum channels ..	404
4.7.1	Projective measurement	245	5.5.4	Gate-based quantum computing	406
4.7.2	Qubits register measurement	248	5.5.5	Distributed quantum computing	410
4.7.3	From computational vector state to full state tomography	249	5.5.6	QEC and FTQC	411
4.7.4	Non-selective and selective mea- surements	250	5.5.7	Quantum parallelism	412
4.7.5	Positive Operator-Valued Mea- surement (POVM)	252	5.5.8	Analog quantum computing	414
4.7.6	Other measurements concepts	252	5.5.9	Quantum thermodynamics	415
4.8	Gate-based quantum computing key takeaways	254	5.5.10	Reality	415
4.9	Bibliography and notes	255	5.5.11	Bibliography and notes	416
5	Quantum computing engineering	259	5.6	Quantum uncertainty	418
5.1	System architecture	259	5.6.1	Optimism	419
5.1.1	Classical computing architecture ...	259	5.6.2	Pessimism	421
5.1.2	Gate-based quantum computers key parameters	259	5.6.3	Middle road	423
5.1.3	Quantum computers segmentation .	264	5.6.4	Bibliography and notes	425
5.1.4	Exotic quantum computing paradigms	267	5.7	Economics	427
5.1.5	Qubit types	269	5.7.1	Pricing	427
			5.7.2	Cost structure	429
			5.7.3	Unit sales and deployments	431
			5.7.4	Benefits	431
			5.7.5	Ecosystem structure	432

5.7.6	Bibliography and notes	435	5.8	Quantum computing engineering key takeaways	436
-------	------------------------------	-----	-----	--	-----

Foreword

The birth of the second quantum revolution may be reasonably settled to 1982, with two celebrated articles: “Experimental Test of Bell’s Inequalities Using Time-Varying Analyzers” by Alain Aspect and colleagues, and the famous theoretical paper “Simulating Physics with Computers” by Richard Feynman. They were quickly followed in 1985 by David Deutsch’s “Quantum theory, the Church–Turing principle and the universal quantum computer”, providing the birth certificate of what is known now as quantum computing. Incidentally, I got a position at CNRS also in 1982, a good spot to witness these developments. However these subjects, including the surprising quantum key distribution scheme proposed by Charlie Bennett and Gilles Brassard in 1984, were still quite confidential at that time. It took ten more years for the topic to attract broad international attention, with the articles “Algorithms for quantum computation: Discrete logarithms and factoring” (1994) and “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”, both by Peter Shor.

This started an explosive development, which has not stopped since then. The big impact of Shor’s algorithm is that it may break cryptographic codes that were considered fully secure at that time, and are still broadly used on internet - so each layman gets concerned. This, and other reasons, triggered a huge international effort to build a quantum computer. Since it was clear that such a computer was still far away, this effort was mostly academic during the next 20 years. But another earthquake happened in 2014, when large US companies, Google and IBM as the most visible ones, decided to get into the race. They set up heavily funded research programs, focused on a specific technology, superconducting quantum bits. This move stimulated a high level of private investment, leading to the creation of a flurry of dedicated companies, initially startups, but then quite respectable enterprises in terms of funding. They were not limited to superconductors, but included all quantum objects likely to become a quantum bit, a qubit : ions, atoms, photons, and even more exotic ones. One must honestly say that, from this time, the language used in this research field changed significantly: the way to speak with investors is not the same as the way to publish articles in reputable scientific journals. The result was a mixture of authentic scientific and technological progress, coming together with hype and overselling. This may be seen either as an unavoidable evil, or as the appropriate way to communicate in an extremely competitive high-tech environment.

A good example is the way to name the “ultimate goal”, the point where a quantum computer can really achieve a useful task that no classical computer can do in any reasonable time. The first name that popped up was “quantum supremacy”, but it did not make it, both for technical and political reasons. The next one was

“quantum advantage”, more politically correct, but still difficult to warrant. The next one is “quantum utility”, clearly more modest, but more likely to be reached in a reasonable time span. In the current language, “utility” is used as a step towards “advantage”, with some possible oscillations between the two due to progress in classical computing.

Then a question arises : given the very fast progress of the devices, their technical complexity, and the strange mixture of science and publicity used by companies to communicate their results, how to find one’s way in the jungle, without getting lost or eaten ? One would need somebody with good scientific and technological knowledge, but also a huge talent for global accounts, given the many facets of these developments, involving physics, hardware, software, middleware, logistics, and finance. We are lucky enough that this person exists, his name is Olivier Ezratty, and you hold his book in your hands or on your computer screen. This is not a modest enterprise, and $\approx 1,822$ pages carefully written in L^AT_EX is a minimal container. Go for it, it’s free ! But how to approach such a monument ? Fortunately there is a reading guide in the first pages, able to orient you depending on your background and interests. It includes also an illustrated “Pandora’s box”, listing the main high level topics (physics, mathematics, human sciences, engineering, computer science, geopolitics), together with their inner content, opening doors towards many stimulating questions for science, technology, and society.

So where do we stand now, already 44 years after the baby was born ? The general idea is that the real stuff, the one needed for supremacy / advantage / utility for practical problems of industrial interest, requires what is called now “Fault Tolerant Quantum Computing” (FTQC), in the line of Deutsch’s ideas; this is quite attractive for computer scientists, but very hard to achieve. Physicists tend to prefer quantum simulators, in line with Feynman’s ideas, potentially useful for more fundamental physics problems. Myself, as a physicist, will go for the second trail, because one should not dismiss fundamental physics problems: it is by solving them that science made progress, and not necessarily by jumping right now into the market. Actually, it is not clear that all fundamental questions to be answered before being able to implement quantum computing have been solved, as summarized in the old joke : “Where is quantum computing? Somewhere between cold fusion and hot fusion”, a physicist’s warning that nature’s laws are not to be fooled around.

Anyway, as we say in French, “*on ne discute pas des goûts et des couleurs*” (there’s no accounting for tastes): whatever your personal interest is - physicist, computer scientist, or investor -, and wherever you may want to look at - computing, communications, sensors or metrology - the development of quantum technologies is a fascinating arena where much more is expected to come.

I cannot close this foreword without a few words about another ongoing revolution, artificial intelligence and large language models generative artificial intelligence (simplified as “LLM”). The most straightforward step forward, taking place currently in all domains of scientific and technical activities, is to use LLM as super-efficient assisting tools; their progress during the 2 or 3 last years has been astonishing, and it cannot be listed

here. Another more speculative idea is to think about a “quantum AI”, an AI running on a quantum computer. The mismatch between the current performances of quantum computers, and what is used now on supercomputers running AI, makes this objective extremely far-fetched. Still, blue-sky thinking about what might be achieved by such machines - and even the question of whether they will still be “machines”, is attracting attention by some quantum theoreticians. Where will these revolutions lead humanity ? Only time will tell.

Philippe Grangier

Director of Research emeritus at CNRS, Institut d’Optique Graduate School (IOGS), and incoming President of the French Optical Society (SFO).

September 2026

Contents

This is **Part 1/5 – Prologue** of the book “Understanding Quantum Technologies 2026”, which is also downloadable using [this link](#).

The downloadable PDFs are available in a single A4 and Letter version, containing the five parts in sequential order. You can also download compressed PDFs for the five parts in A4 and Letter formats. Their size fits the constraints of ebook readers like the Kindle from Amazon.

The e-reader and printed version of the book separates parts 1, 2, 3, 4 and 5.

- The **first part** contains a history of quantum physics, some quantum physics 101 and everything about quantum computing basics and engineering including quantum error correction, the energetics of quantum computing, a view on pessimism and optimism with regard to the advent of scalable and useful quantum computers, as well as some economic considerations.
- This **second part** contains the sections dedicated to quantum computing hardware with all types of qubits and their related industry vendors, enabling technologies, and unconventional computing solutions which are potential alternate routes between classical and quantum computing.
- The **third part** contains the sections dedicated to algorithms, software tools, and applications across multiple markets. It also contains an inventory of quantum computing software and service vendors.
- The **fourth part** covers quantum communications, quantum cryptography, post-quantum cryptography and quantum sensing.
- The **fifth part** describes the quantum ecosystem at the entrepreneurial level, the country level, on various societal topics, corporate adoption methodologies, quantum pseudosciences. It ends with a glossary, and book bibliography.

The book is split into five parts to make its printing easier, some online printing services including Amazon being limited to a maximum of 600 pages.

You can order the printed version of this book in five parts on all Amazon sites by searching for the book title, edition 2026, when it is released around October-November 2026.

1 Why

Welcome to the 9th edition of “Understanding Quantum Technologies”. This book remains your Hitchhiker’s guide to the quantum galaxy. It keeps being unique in many respects: its breadth, its density, and its purpose. It presents a kaleidoscope for quantum domains with a 360° perspective encompassing historical, scientific, technological, engineering, entrepreneurial, geopolitical, philosophical, and societal dimensions. It is not a quantum for dummies, babies, or your mother-in-law book. It mainly targets information technologies (IT) specialists and engineers who want to understand what quantum science and technologies are about and decipher its ambient buzz, but also all participants in the quantum ecosystem from academics to industry vendors, investors, and policy makers, and at last students who would like to investigate quantum science and technologies as an exploratory field. I have updated the book every year like a Swiss clock every September since 2018.

While being rather scientific and technical in many parts, it tries to explain things and translate the complex quantum lingua in other tech’s lingua, particularly for IT and computer science professionals. It looks at the history of quantum science and ideas and pays tribute to key people, from the past and the present.

After providing the basics on quantum physics and quantum information science, the book investigates rarely covered aspects of quantum technologies and engineering like various enabling technologies (cryogenics, electronics, materials design, semiconductors, cabling and lasers, manufacturing techniques), their energetic dimension, what raw materials are used and where they come from. It even explains how research works in general and in the quantum realm and its various codes and practices, nowadays being revolutionized by generative AI and LLMs. The book can also be viewed as an integrated collection of several books, which also covers quantum sensing, telecommunications, and cryptography. It also contains many crafted custom illustrations that I use in my teachings, training and other talks.

The book also calls out the nonsense when necessary. It is abundant, particularly when some analysts and vendors are fueling the quantum hype. As quantum technologies are more commonplace, these are still largely misunderstood by general audiences as well as by many IT professionals, and by many people writing about it. I’ve been through many technology waves in the last four decades, but this one is particular with a wider gap between the hype and the reality.

Governments’ technology sovereignty ambitions and industry vendors hunger for funding elevated quantum technologies to the rank of strategic sectors in many developed countries, on par with artificial intelligence. Most governments have launched their national quantum plans, starting with Singapore, the UK, China, USA, Germany, Japan, Australia, France, Russia, Israel, Taiwan, India and the Netherlands. The worldwide quantum technologies race is on. Countries are embattled to acquire or preserve their technological independence or leadership, like if it were the last chance to achieve it, particularly for those countries that felt they

had lost the digital battle against the USA and Asia (mostly China, South Korea and Taiwan).

Like many deep techs, quantum technologies are dual-use ones, with both civilian and military use cases, increasing their strategic stakes.

While it has not yet reached the volume and funding of other sectors such as artificial intelligence or the digital cloud, the quantum startup and small business ecosystem continues to expand worldwide. In this book edition, I mention about 970 such companies in many different categories (hardware, software, telecommunications, cryptography, sensing, enabling technologies, services), +100 vs 2025. In most cases, hardware is in the deep tech realm if not in hard tech territory, with many still at an applied research stage with a rather low technology readiness level. Most companies in this sector are R&D labs on wheels with relatively low technology readiness level offerings.

The force of the hype is strong in quantum. A few scientists, their laboratory’s communication department, startups and large vendors frequently oversell the impact of their work. The hype shows up when some are pretending that quantum computing is ready for business and misleading customers about the maturity of the technology. And it is hard to detect it without being specialized enough.

This book has another special flavor. It is the result of a human adventure at the heart of the quantum ecosystem. Back in 2016, I decided to select the theme of quantum computing for my usual techno-screening activities, ranging from preparing conferences and training to writing educational ebooks for professionals. With my friend **Fanny Bouton**, we prepared and ran a popularization conference on quantum computing in Nantes, **Le quantique, c’est fantastique** in June 2018 [1] and did numerous subsequent presentations. She brought and still brings a different perspective, and now a clear quantum business acumen as the head of quantum at OVHcloud.

This book started as a compilation of 18 posts published in French between June and September 2018. It was an

extended transcript of the above mentioned conference. After two enriched editions in French in 2019 and 2020, I switched to English in the fourth edition in September 2021, and its subsequent editions until this one.

Some of this book’s readers call it a bible or an encyclopedia. It is more like a Yahoo! of quantum for those old enough to remember it. It contains over 11,700 bibliographical entries, mostly made of scientific papers, making the “bibtex” L^AT_EX bibliography file probably one of the largest ever assembled. It also contains about 1,300 illustrations.

In this journey that is still going on, I had the opportunity to meet with top researchers and entrepreneurs, first in France, and then internationally. This list keeps growing. I thank them all in the conclusion of this book. In short, during these years, we have been “embedded” in the scientific and entrepreneurial ecosystem. We also applied one of Heisenberg’s principles derivatives, namely that a measurement device may influence the measured quantity. It was and remains a beautiful adventure with real people, passions, convictions, ups and downs, and, in the end, a nice result with French and European research and entrepreneurship in quantum technologies that are more dynamic and better positioned than a few years ago.

Q How is this book made?

This 9th edition is the second edition of this book written in L^AT_EX. I switched from Word to L^AT_EX to prepare it for the 8th edition in 2025. Before this, I was using Word to write all my books. Translating this book from French to English in 2021 happened to be much easier than the migration to L^AT_EX. It took me about two months to do it, with a mix of automation with various scripts and regex, complicated images conversions or recreation and manual editing. My L^AT_EX preamble has now about 2,150 lines! I tried to follow the rules of scientific-paper publishing. You can recognize a template familiar in arXiv and PRX Quantum. I use mainly the TexStudio editor. Fortunately, I got the help from LLM-based chatbots with a mix of ChatGPT, CoPilot, Google Gemini and now, Claude Pro. I describe my whole migration journey in the “Making of” section in the appendix of the 8th edition of the book. In this version, this appendix is entirely focused on the way I have used LLMs to improve this book and even create specific content. In a way, it’s a giant “AI disclaimer” for the book.

You may wonder how this book can be free. I have published all my books like this since 2006 on entrepreneurship, artificial intelligence and other technology, and science related topics. I have fared well since then. I

favor distribution breadth over direct revenue. It makes knowledge easily accessible to broad audiences, particularly with students. Thanks to a digital format distribution, I correct and update my books on a regular basis. It is quite practical since I mention hundreds of people and organizations, and deal with complicated scientific matters. Afterwards, I sell my time in a rather traditional way with speaking, teaching, training, expertise, and a few consulting missions. I also undertake pro-bono activities in research and as a cofounder of the Quantum Energy Initiative. The business model is simple: the (very) long version of the book is free and the (too) short versions are charged. Since the people who don’t have time usually have money and the other way around, it works quite well even if it may be counterintuitive in the first place. My pride comes with meeting young professionals in the quantum ecosystem who thank me for the book, which contributed to their interest in quantum and starting a research career or working in the industry.

1.1 A domain in search of pedagogy

After having swept through many areas of science and deep techs, I can definitively position quantum physics and quantum computing at the complexity scale apex. Quantum physics is difficult to apprehend since relying on counter-intuitive phenomena like wave-particle duality and entanglement, and on a mathematical formalism that is not obvious to most people, particularly with most IT specialists and developers, one of the key audiences for this book. It is still an open challenge to first understand, then translate this scientific field lingua into natural language for most people, even with a strong engineering background. And you know, I also don’t have a PhD in quantum physics (private joke)!

You probably heard about the famously rehashed quote from **Richard Feynman** who pointed out that when you study quantum physics, if you think you understood everything, you are making a fool of yourself. **Alain Aspect** expressed doubts about his own physical understanding of quantum entanglement and nonlocality that he tested in his famous 1982 experiment which led him to be awarded the Nobel prize in physics in 2022 [2].

Explaining quantum technologies is thus a new and difficult art, but not an impossible one [3]. At least, it is a very diverse one with many different approaches. When reading quantum physics books, you discover a mathematical formalism and many terms like observables, degeneracy, gentle measurement, unitary and projector, operators and the like, and wonder how they relate to the physical world. Sometimes, it takes quite a while before being able to make the connection, whenever possible! On the other hand, you hear simplistic descriptions of quantum physics, particularly on superposition and entanglement, and quantum computing, some coming from quantum computing vendors themselves [4]. And you have the frequently oversimplified Schrödinger-cat

thought experiment, which was devised to expose the measurement problem but is still not well interpreted.

Once you think you understand it after having created a mental view of how it works, your explanations become quickly inaccessible to non-experts. How do you avoid this side effect? Probably by finding analogies and using more visual tools to explain things than too much mathematics. I test this approach in many sections of this book, but mathematics are always useful in many parts. Also, to make sure it does not lose its scientific soundness in the process, many parts of this book past editions have been fact-checked and proofread by quantum scientists. I would say, it's still not enough. You'll be the judge.

This book frequently responds to questions like what, why, where and how? Has Moore's empirical law really stalled?

What does being "quantum" mean for a product or technology? Do we really have objects sitting simultaneously at two different locations? Why parallel opposite vectors in the Bloch sphere representing a qubit state are mathematically orthogonal? Why and where density matrices are useful? What are pure and mixed states describing in the physical world? Why superposition and entanglement are the two sides of the same coin?

When will we have a "real" quantum computer? How can you compare such and such quantum computer technology and qubit type? What is your preferred one (none)? Which one can scale best (all have limitations)? Why do we need to cool many qubit systems at very low temperatures? How are cryostats operating? What is the energy consumption of a quantum computer?

How much data sits in a quantum register? How is data loaded in a quantum program? What data is generated by quantum algorithms and how is it decoded? Are quantum computers made for big data applications? Can NISQ bring some commercial value? Can analog quantum computing compete with gate-based models? Is the Shor integer factoring algorithm a serious threat for your cybersecurity? Will quantum computers save the world (healthcare, climate change, ...)? Have we really achieved quantum supremacy? What is the difference between quantum supremacy, advantage and utility? What is the real speedup of quantum algorithms? How to analyze a quantum computing case study? Are there quantum computing case studies in production?

Will a quantum Internet replace the existing Internet? Can quantum telecommunications enable either faster than light communications or high-throughput data links? How are classical computing technologies competing with quantum computers? Why are some quantum random number generators not that random? Why can entanglement improve the precision of quantum sensors?

Is China going to kill us (metaphorically) with their (not so) huge R&D investments in quantum technologies?

Have they really invested \$15B in quantum technologies? Is Europe doomed against the USA and China?

Oh, and if I'm in an organization... what should I do? Am I late in the game when doing nothing? Should I stay or should I go? Why are some people overselling the capabilities of quantum computers? Will governments build dangerous weapons with quantum technologies?

But you know what? Some parts in this document contain stuff that I write but do not understand well. Or sometimes, I understand it well but when I review it later, my understanding is gone (like for the topics on the right of the complexity scale in Figure 1). Quantum scientists sometimes feel the same.

So, you may understand why I am kind of annoyed when I am invited to present the whole field of quantum technologies in a half-hour session!

1.2 A new technology wave

Quantum computing stays on top of the various applications of the second quantum revolution. Quantum sensing is more exotic and fragmented, and quantum telecommunications and cryptography are less fascinating. Why is quantum computing becoming an important topic? Firstly, because large IT companies such as IBM, Google, Intel and Microsoft are making headlines with impressive announcements that we must, however, take with a grain of salt, with a lot of hindsight, and decipher calmly. Shor's algorithm creates a real long-term threat to RSA and elliptic-curves based cryptography, although the arrival date and resources required for a cryptographically relevant quantum computer remain uncertain.

Above all, it is linked to the broad impact that quantum technologies could have on many scientific fields and digital markets. It may theoretically make it possible to solve problems belonging to classes of complexity that even the largest giant supercomputers will never be able to tackle. The other reason for this sudden interest is that we are still at the beginning of the story. New leaders will show up. A new ecosystem is being built. This is a field where there are still enormous scientific and technological challenges to overcome. It is a land of opportunities for science, technology, and innovation. Like with quantum physics, we are in a highly indeterministic world. It is quite difficult to evaluate the feasibility of large-scale quantum computing. For most scientists, we are still between one and two decades away from it. Some believe it will never show up. Others are more optimistic. One key uncertainty is about our ability to control large sets of entangled qubits with good fidelities. The plan is to fix that with quantum error correction and logical qubits made of physical qubits. It then becomes, at least, a physical scalability issue with a bunch of complex engineering issues related to cooling, cryo-electronics, cabling, classical computing, miniaturization and interconnectivity.




Figure 1: A scale of complexity in quantum physics and technologies, from the easy (left) to the very difficult (right), at least, as far as I am concerned. I could have added qRAM here but there was not enough room in the chart! © Olivier Ezratty, 2023-September 2026.

It is a very interesting living case study of how mankind builds upon scientific progress and addresses the most difficult challenges around. In this respect, it is on par with controlling nuclear fusion. The joke being, who is going to be first? Nobody really knows for sure.

1.3 Reading guide

Here is an attempt to prioritize which parts of this book you could read according to your business and scientific level (Figure 2). Physicists can find a state-of-the-art tour covering all dimensions of quantum technologies beyond the field they have already mastered.

Computer scientists, engineers and students in various scientific fields are the core target audience for this book, as it presents, popularizes and contextualizes the various scientific, mathematical and engineering concepts used in quantum technologies.

The required level of mathematical and computer basics is at the bachelor's degree level for most parts and sometimes bachelor or master levels. Non-technical and decision-makers can still read the sections dealing with usages as well as how countries are faring and societal issues.

Figure 3 shows another view of the table of contents, with the overall logic between the lower « physics » layers and the upper hardware, software and solutions layers, with the parts in yellow being the most difficult to read for a non-scientific audience, although it may also be the case in other parts.

Let's also mention one of the reasons why a curious mind may like quantum technologies: it encourages you to explore many scientific disciplines, even human and social sciences, like a scientific Pandora's box as shown in Figure 4. On top of that, learning quantum science is probably more efficient than Sudoku or crosswords to train your brain muscle as it ages!

Q How quickly can this book be outdated?

This book mixes three content types with different shelf life. History, quantum physics and mathematics is relatively long lived.

Algorithmic and error-correction state of the art changes regularly. Vendor roadmaps, funding rounds, headcounts, and company existence change even faster. I cared in this edition about dating events and papers whenever possible to make them well positioned in time.

However, as I update the book every September and until December, it's relatively up to date, at least compared to the current book literature from traditional publishers. Most books are never updated and become outdated quite rapidly.

When company are mentioned with their logos, the format includes their country of origin, year of creation, and approximate total funding. This aggregates seed

book sections		quantum physicists	computer scientists and developers	students in sciences (STEM)	non IT audiences	policy makers
volume 1	why				↓	
	history and scientists					
	quantum physics 101	known	optional			
	gate-based quantum computing					
	quantum computing engineering					
volume 2	quantum computing hardware					
	quantum enabling technologies		optional			
	unconventional computing					
volume 3	quantum algorithms					
	quantum software development tools					
	quantum computing applications					
vol 4	quantum communications and cryptography					
	quantum sensing					
volume 5	quantum ecosystems around the world					
	corporate adoption					
	quantum technologies and society					
	quantum fake sciences					

Figure 2: Understanding Quantum Technologies parts and audience relevance. Color codes: dark green (highly appropriate), light green (optional), grey (probably irrelevant or too technical), orange (not relevant since really too technical). © Olivier Ezratty 2021-Septembre 2026.




Figure 3: How the topics covered in Understanding Quantum Technologies are related with each other. The orange parts are the harder to read for the quantum and scientific non-expert, then the blue parts which are still scientific but somewhat easier, and the green part which is easier to access to general audiences. © Olivier Ezratty. 2021-September 2026.

and private equity funding, IPO and SPAC related cash funding, and sometimes various forms of public funding depending on the way companies are accounting for their funding source. This provides some raw indication of their funding level but can change over time. Some

trusted sources can be used for checking the data like the Crunchbase database.




Figure 4: The many scientific and other domains to explore when being interested in quantum technologies. That's why you will like this book if you are a curious person. © Olivier Ezratty, 2021-2026.

1.4 First and second quantum revolutions applications

Quantum physics has been implemented since the post-war period in almost all products and technologies in electronics, computing, and telecommunications.

This corresponds to the **first quantum revolution**. It includes transistors, whose first point-contact version was demonstrated in 1947 by John Bardeen and Walter Brattain [5]. That first device was a bipolar transistor. The MOSFET which uses the field effect and is the basis of our existing digital world came later, in 1959, with Mohamed Atalla and Dawon Kahng at Bell Labs. The first quantum revolution also covers photovoltaic cells which rely on the electron-hole pairs created by incident photons, and lasers which also exploit the interaction of light and matter and are used in a very large number of applications, particularly in telecommunications and optical storage (audio CD, DVD and the like, which are now mostly outdated). Many medical imaging solutions rely on various quantum effects, including nuclear magnetic resonance imaging (MRI). LEDs are also based on quantum effects. GPS relies on atomic clocks synchronization. Quantum dots used in high-end LCD displays and Smart TVs also use variations of the photoelectric effect. Alexey Ekimov (1945, Russian), Louis Brus (1943, American), and Moungi Bawendi (1961, American-Tunisian-French) were awarded the Nobel prize in chemistry for the discovery and synthesis of quantum dots in October 2023. The list is long, and we will not detail all these use cases!

The **second quantum revolution** covers the technologies combining all or part of the ability to control individual quantum objects (atoms, electrons, photons), use quantum superposition and entanglement. The notions of first and second quantum revolutions were created in 2003 independently and almost simultaneously by

Alain Aspect on one side and by Jonathan Dowling and Gerard Milburn on the other side. In the United States, the paternity is attributed to the latter two, while in France, it is attributed to the former! A “déjà vu”. The expression appeared in the preface by Alain Aspect of a book by John S. Bell in June 2004 which was written in February 2003 [6]. Jonathan P. Dowling and Gerard J. Milburn also coined the expression in a June 2003 paper [7]. Later, Dowling made a large inventory of various quantum technologies embedded in this second quantum revolution [8, 9].

The scope of the second quantum revolution covers various recent applications of quantum physics that integrate quantum computing, quantum communications, quantum cryptography, quantum sensing and quantum matter. Said simply, it is about improving our digital world performance and security, and to increase the precision of all sorts of sensors and in the case of quantum matter, to invent new materials (Figure 5).

- **Quantum computing** is the broad domain of using quantum physics to find solutions to various computing problems. It includes various computing paradigms like gate-based computing, quantum annealing and quantum simulations. Hundreds of pages cover this topic in this book, from hardware to software.
- **Quantum cryptography** is a means of creating shared secret symmetric encryption keys between distant parties thanks to quantum physics phenomena and rules, like photon entanglement and the no-cloning theorem. It relies either on fiber optic communications or on space links with satellites as China has tested with its Micius satellite since 2017. Even though some researchers are proposing to use new quantum computer cryptography schemes, most




Figure 5: The second quantum revolution applications. Quantum computing includes analog and digital or gate-based quantum computing. Quantum cryptography is one application of quantum communications, which also include quantum computer and quantum sensor quantum interconnect. Quantum matter is not yet commonly accepted as being part of the second quantum revolution. © Olivier Ezratty, September 2026.

quantum cryptography plans rely on using quantum key distribution using photonic links.

- **Quantum telecommunications** enable distributed computing, connecting quantum computers enabling qubit to qubit distant entanglement, and, potentially, quantum sensors, which can be implemented to improve their accuracy. This field is still in the making and could become the base for a very secure quantum Internet and quantum cloud infrastructures. We cannot exploit it to transmit classic information faster than today. However, it can be used to distribute quantum processing on several quantum processors. It could provide a mean to “scale out” quantum computers when “scale-in” approaches reach their limits. This requires a lot of engineering, particularly to convert solid qubits into photon qubits deterministically and leverage shared entanglement resources.
- **Post-quantum cryptography** is a different field which is intended to replace current classical cryptographic solutions with new solutions that are supposed to be resistant to attacks carried out by future quantum computers. It does not belong to the second quantum revolution per se but is rather a consequence of it.
- **Quantum sensing** makes it possible to measure most physical dimensions with several orders of magnitude better precision than existing classical sensing technologies, even existing atomic clocks. It is a vast scientific field that is the subject of numerous research projects and industrial solutions. It includes ultra-precise atomic clocks, cold atom accelerometers and gyroscopes that use atomic interferometry, SQUIDs (superconducting based) and NV

center magnetometers. Micro gravimeters measure gravity with extreme precision, enabling discoveries of underground anomalies like holes, water, and various materials. This domain also includes various advanced medical imaging systems with higher precision and non-destructive imaging and measurement tools [10, 11]. A dedicated section of this book covers quantum sensing with its broad scope of solutions or prospect solutions.

1.5 Why quantum computing?

Quantum computers are bound to solve complex problems that are and will stay inaccessible to classical computers. This happens when these problems’ solutions scale exponentially in computing time on classical machines. In extreme cases, computing time on conventional computers for exponential problems, even with the most powerful supercomputers of the moment, could largely exceed the age of the Universe for practical problems, estimated at 13.8 billion years. We cannot be that patient!

The ambition for quantum computers it to solve these intractable problems in times that scale differently, polynomially, or even at a lower scale (linearly, logarithmically, ...) and, preferably, down to reasonable times depending on the business needs. Obtaining some practical speedup with a quadratic or polynomial gain over a classical exponential problem seems not obvious for multiple reasons, the main one being the relatively slow quantum gate operations. Otherwise, quantum algorithms may under certain circumstances bring other benefits like a better solution quality. A big disclaimer




Figure 6: Simplified view of the quantum computing theoretical promise. The chart on the left shows the simplified view of how computing time scales exponentially on a classical system with a computing time T_c , corresponding to a lower computing time T_q for a quantum solution, presented with a computing time log scale on Y and linear problem size scale on X. The practical question for a given problem and system is to identify with which problem size T_q gets smaller than T_c . The chart on the right, inspired by [12], zooms on this crossover zone, using a X qubit scale which is an approximate scale for the problem size. It shows in a simplistic manner that the crossover T_p with a polynomial speedup happens much later than with an exponential or superpolynomial speedup with T_e . This crossover time must be acceptable from the user standpoint. It depends on many factors including the qubit technology used in the quantum computer. As we will see later in the book, this time may not be acceptable for the quantum algorithms which bring only a polynomial or quadratic speedup, like many optimization related algorithms using Grover's algorithm. It is frequently evaluated in years and even centuries. This is due to the prefactors including slow quantum gates, and potentially, data preparation and loading initial costs. This is why we need not only algorithms bringing some theoretical superpolynomial or exponential speedups, but solutions with a practical speedup for real world problems with real-world data, and considering all their computing costs. Before delivering this promise, quantum computers may bring other benefits like producing better and more accurate results and/or doing this with a smaller energy footprint. © Olivier Ezratty, 2022-2026.

is that all these benefits are at this stage an undelivered *promise*, at least for useful commercial applications. Turning these promises into reality is one of the most difficult, challenging and exciting goals in science and technology development.

1.5.1 Quantum computing promise

Typical exponential problems are combinatorial optimization searches and quantum chemical simulations. Their size is usually expressed as a quantity like a number of locations with a traveling salesperson problem (TSP, with various constraints) or a number of molecular orbitals for a chemical simulation. Exponential problems are “intractable” because their classical computation time evolves in crazy proportions, exponentially, with their size.

Quantum physics and chemistry simulations come first, governed by many-body quantum physics equations. Typical chemical quantum simulations algorithms determine the minimum energy configuration of a system, its ground state, with a greater precision than classical methods. Other quantum algorithms could help determine how molecules interact in chemical pathways and at molecular dynamics. This will not go so

far as simulating an entire living being or even a single cell. It will already be a fantastic feat to simulate some simple de-novo protein folding in a better way than what DeepMind AlphaFold 3 is doing today using deep learning techniques, the next step being protein interactions simulations. The competition from classical machine learning is still significant and growing [13]. Physics simulations also deal with material designs based on the understanding of crystal structures or how magnetism operates.

Optimization problems are also in sight, and you can find them in many industries, like in transportation, logistics, retail, manufacturing and financial services. One such problem is the job scheduling problem, on how to optimize all human and physical resources in a factory. There are however, few algorithms bringing a clear exponential speedup here, and a polynomial speedup often shows up practically only for very large problem size and long computing times, when compared with classical solutions. There, you have to learn about comparing classical and quantum heuristic solutions.

Engineering problems in demand of partial differential equations solutions is another broad field of interest, particularly in aerospace engineering and fluid mechanics domains.

Quantum machine learning is another domain of interest with training and inferences in machine learning and neural networks models. It is within the reach of conventional computers equipped with GPGPUs (general purpose GPUs) such as Nvidia’s H200 and B200 and their tensor processing specialized units which are optimizing matrices operations. Obtaining a quantum advantage is less obvious in this field, particularly since machine learning must usually be trained with large classical data sets. These large data sets must be prepared classically beforehand. Their loading in a quantum computer is generally slower than on a classical computer. It can remove a potential quantum computing speedup unless an efficient state-preparation or data access model is available. However, quantum machine learning may potentially bring some other benefits like the creation of better solutions in terms of prediction accuracy, instead of bringing some speedup.

“Building a quantum computer is a race between humans and nature, not between countries”.

Lu Chaoyang, China, December 2020.

Cryptanalysis comes last, which is of particular interest to the NSA and their peers to break RSA-type public-key encryption security. There is no business case for this unless you want to spy on somebody. Or you want to go to prison, depending on the country where you live.

Solving these problems with quantum computers is still a quest. It depends on many factors and unknowns: the creation of more efficient quantum algorithms, the evolution of methods for preparing and encoding data in quantum registers, improving the quality of qubits and at scale, the speed of execution of quantum gates, the creation of even more efficient error correction codes, the ability to quantum-interconnect quantum processors, and in some cases, the ability to parallelize quantum circuits on a large number of machines, at an affordable cost, to reduce computation times. The diversity of these challenges, the solutions envisioned by academics and industry vendors, and the associated scientific, technological and economic uncertainties make any prognosis hazardous. One thing is almost certain: given the cycles of development and experimentation, even with being optimistic, solving all these problems will take a long time.

Also, quantum computing will not become a “jack-of-all-trades” solution nor a replacement tool, but more a complement to current High-Performance Computers (HPC). Many, if not most of today’s classical computing problems and software are not relevant use cases for quantum computing. Most business data processing tasks will remain classical, like running most ERPs (enterprise resources planning), accounting, email, databases, media applications, and others.


From an economic historical perspective, the consequence is that quantum computing will probably not be a Schumpeterian innovation engendering a “creative destruction” process. Quantum computing is more likely to complement classical HPC than replace general-purpose computing. It is an incremental instead of being a replacement technology even though what it may achieve is beyond reach as of today. As a result, there may not be much value migration happening within the computing sector [14].

You probably will not have a quantum desktop, laptop, or smartphone to run your usual digital tasks although quantum technologies can be embedded in these devices like quantum sensors and quantum random number generators. Quantum computers will be hidden from users and sit in cloud data centers, like Nvidia GPGPU racks.

This will be even amplified by the progress we can anticipate with wireless telecoms. When and if quantum computers scale, sometime after 2030, we will probably rely on 6G or 7G networks with even better latency and bandwidth. Of course, it is still hard to anticipate the usages brought by quantum computers when they will scale. Let’s also boil in the fact that, as we will see later, quantum computers are not excellent at handling big data. They will also probably be generally ill suited to low-latency, high-throughput real-time workloads, particularly with significant data movement. This makes it less relevant to use a local quantum processor, as it makes sense today to have local neural networks capacities to handle your in-camera image recognition processing and voice recognition in smartphones. Less data means more relevance for distant quantum computation implemented in the cloud.

Business cases are investigated for different markets such as retail, transportation, logistics, telecommunications, financial services and insurance. Many businesses have complex optimization problems to solve. Like with most technology-driven disruptions, businesses may progressively discover quantum computing use cases as its market and related skills grow.

We will avoid putting the cart before the wheel. Contrarily to what is usually said, we do not lack algorithms and use cases. What is missing the most is the hardware to run it. All these promises are dependent on the ability to create large scale and fault-tolerant quantum computers, which are years if not decades away. In the interim, we may end up having quantum systems able to deliver other benefits like producing better and more accurate results and/or doing this with a potentially smaller energy footprint, but not with some exponential computing time speedups. They can also help us learn quantum computing and, indirectly, improve legacy classical solutions. However, what we lack are quantum algorithms with reasonable demonstrated, theoretical and/or practical quantum advantage versus the best classical methods.



empirical observation from 1965

the complexity of integrated circuits doubles every 18 months

many derivatives with :

- transistors density
- cost / transistor
- supercomputing power
- storage capacity
- cost of storage / GB
- networking speed
- CMOS imaging sensors resolution
- human genome sequencing cost





Fig. 2 Number of components per integrated function for minimum cost per component extrapolated vs. time.



also, an economics driven law

Figure 7: Gordon Moore's original 1965 paper dealt with both transistor number per chip trends and an economic driven law. Beware, the "doubles every 18 months" mention in the assembled chart is the later David House variant, the 1965 paper predicted a doubling every year. Key drivers of this law were probably the emergence of the personal computer market in the late 1970s and then, the smartphone and the consumerization of all digital tools. It drove the creation of mass market which fueled Moore's law and its many derivatives on storage and telecommunications. Source: [15] © assembled by Olivier Ezratty, 2023.

1.5.2 Moore's law

One strong motivation to build quantum computers is the perception that classical technology progress may be stalling. Moore's law effects are supposed to end. Classical computing progress seems to have reached hard limits, and a disruptive approach is needed. That may not be true.

Gordon Moore's law was a sort of exponential regression used to predict the rate of growth of the number of components in a chip. His 1965 paper actually predicted a doubling every single year over the following decade, from about 60 components in 1965 to 65,000 in 1975 [15]. He recalibrated it to a doubling every two years in 1975 [16]. The popular "18 months" variant came later from his Intel colleague David House and applied to computing performance, not to component count. Gordon Moore's paper was written when he was working at Fairchild Semiconductor, only 5 years after the production of the first integrated circuit and 6 years before Intel created its first microprocessor, the 4004. It was an era of relatively fast technological progress. The Moore's law nickname was coined about a decade later by Carver Mead, a Professor at Caltech and friend of Gordon Moore. Moore passed away in March 2023 while Mead is still with us. Moore's law was based on a sampling made with only five data points ranging from 1959 to 1965 as shown in Figure 7, in the very early years of the history of integrated circuits production.

Integrated circuits were invented twice and independently, by Jack Kilby at Texas Instruments in September 1958 with a germanium hybrid circuit, and by Robert Noyce at Fairchild in 1959 with a silicon planar version which is the one the industry built upon. First commercial production started in 1960-1961. The progress was both in number of transistors, cost per transistor and surface density. A regular wafer was only one inch large when nowadays, they are 12 inches large (30 cm) and can accommodate hundreds if not thousands of chips depending on their size, or just one large chip, like Cerebras' giant WSE-3 wafer-scale chip manufactured by TSMC and packaged in their CS-3 server.

Gordon Moore's empirical law application would have a marginal impact on computing times for exponential problems. Whatever the progress, it would not bring the capacity to solve exponential problems in non-exponential times. The extension of Moore's law with "More than Moore" architectures won't change this. Instead of just shrinking transistor size of stacking them, "More than Moore" focuses on functional diversification of chips with system integration with combining digital and non-digital components like sensors, RF and power management into a single chip, which is seen in embedded systems like smartphones or various connected objects. It also involves the use of new materials to enhance performance, heterogeneous integration with stacking and packaging different types of chips together on a chiplet substrate, and enabling chips to interact

with the physical world with features like sensing, actuating, and communication.

Adding one qubit doubles the dimension of the system's Hilbert space handled during a quantum circuit execution. This exponential state-space growth is a resource for some quantum algorithms, but it does not automatically double quantum computing power. Also, maintaining qubit quality as systems scale is increasingly difficult due to noise, and particularly with the overhead of quantum error-correction when supporting large circuits.

So, why does Moore's law seem to have reached its limits? As a matter of fact, it hasn't yet, when looking at the trend plotted in Figure 8. The number of transistors per chip is still increasing. The progress that is not literally associated with Moore's law and that has stalled is elsewhere, with the single-thread performance, chip clock rates, power per chip and their number of logical cores.F

Dennard scaling is the real law that came to an end around 2006. It stalled semiconductor progress in the three mentioned areas (thread performance, clock, power). Robert Dennard (1932-2024, American) and his IBM colleagues published this scaling rule in 1974 [17]. He forecasted that, as transistors density increases, the power consumed per unit area of the chips would be stable. As shown in Figure 9, this happened since the transistor's voltage and current could decrease with their density, while increasing the clock frequency. Starting with 65 nm integration in 2006, this rule came to an end, as supply and threshold voltages ceased scaling proportionally, while leakage and power-density constraints became dominant. Another phenomenon is the tunnel effect happening at the thin grid oxide level, that is reduced using high-dielectric constant oxides ("high k dielectric"), introduced by Intel at the 45 nm node in 2007.

That's why, among other phenomena, your laptop computer is also heating your legs when you use it in public transportation or in your coach. As a result, this "heat barrier" limits the capacity to sustain processor clock speeds much beyond 5 GHz with conventional air cooling [18]. Stock high-end desktop parts now reach a bit more in single-core turbo mode, like the Intel Core i9-14900KS and its 6.2 GHz clock rate. Extreme overclocking goes further, with a 9.2 GHz world record set in 2026 on an Intel Core i9-14900KF cooled with liquid helium, which is obviously not very practical [19].

The transistors current leaks started to grow and power consumption soared. This is what prevents the growth of processors clock. At the beginning of the 2000s, Intel planned in its roadmaps to raise their CPU clock frequency up to 20 GHz.

Intel then stopped playing this game and instead entered the multicore realm. However, in June 2021, Intel released a new microprocessor for high-end laptops running at a 2.9 GHz base clock but with a 5 GHz turbo

mode for a single core, the 4-core i7-1195G7, etched in 10 nm, and with a 28W TDP (thermal dissipation power).

The semiconductor demand switched in 2007 towards low-power multi-functions chips for smartphones. This opened a boulevard for Arm core-based processors and growth for corporations like Qualcomm.

Koomey's law empirical law proposed in 2010 by Jonathan Koomey observed that the available computing power per consumed kW increased steadily, doubling every 1.57 years between 1946 and 2009 [21]. However, this doubling slowed down to 2.6 years after 2000, due to the end of Dennard's scale. It indirectly explained why multicore architectures are limited in number of independent cores.

Multi-core architectures enabled deeper and larger parallel processing but still within limits previously formalized by Amdahl's law in 1967 [22], which describes the upper limits of parallel computing systems acceleration, and at the expense of increased technological and programming complexity which adds one more level of parallelism.

Dark silicon is a phenomenon associated with the end of Dennard's scale. As a chip gets too hot, it becomes difficult to use it entirely. Various methods are then combined to circumvent this inconvenience: on-demand cores or functions deactivations according to usage needs, a shutdown of certain portions or cores, a voltage drop, a selective clock frequency adjustment per core or simply, a low clock speed (Nvidia GPGPU's run at about 1 GHz). This is used in the Arm core-based processors of smartphone chips, whose cores do not use the same clock rates, in the so-called big.LITTLE architectures created in 2011, and replaced with the more flexible DynamIQ architecture in 2017. There are many other techniques to improve classical processors energy efficiency [23].

Some other laws are also applicable in the science-fiction domain when you reach quantum limits [24].

1.5.3 Transistor density evolution

The semiconductor industry had to cope with many limitations when improving transistor density, Landauer limit, the heat barrier, some unwanted quantum effects, the reticle size limits and the resolution of etching manufacturing techniques.

Landauer barrier defines the minimum energy required to erase a bit of information, $k_B T \ln 2$, which amounts to about 2.8×10^{-21} J or 0.018 eV at 300 K, and to about 10^{-25} J at 10 mK in a dilution refrigerator. It is a very low theoretical barrier contested by some physicists, although it was verified experimentally in 2012 with a single colloidal particle trapped in a modulated double-well potential [25]. And it can be circumvented as we will see with the technique of adiabatic and reversible computing. It was defined by Rolf Landauer (1927-1999, researcher at IBM) in 1961.




Figure 8: 50 years of microprocessor technology trends. It shows that transistor counts continued to rise over the plotted period per chips (or chiplets), although single thread performance, clock, power and number of logical core have stalled since about 2005 due to various empirical laws limitations like Dennard's scaling, Amdahl's law which describes the upper limits of parallel computing systems acceleration, and Koomey's law end, according to which the available computing power per consumed kW increased steadily, doubling every 1.57 year. Source: [20] with data updated in 2022, and additions © Olivier Ezratty, 2021-September 2026.




Figure 9: Dennard's scaling which explains the dark silicon phenomenon where all CMOS chips components cannot be used simultaneously at optimum performance. Compilation © Olivier Ezratty. 2020-2023.

Quantum effects such as unwanted tunneling can degrade conventional CMOS behavior at very small feature sizes, although tunneling is deliberately exploited in other device classes.

Etching resolution is getting smaller to enable the manufacturing of more precise and smaller features in transistors, particularly below 10 nm nodes. Lithography etching systems are using extreme ultraviolet, coming from ASML. Etching resolution indeed depends on the wavelength of the light used to project a mask on a photoresist. Lowering the transistors size requires increasing this frequency to decrease the wavelength,

and thus go from the current deep ultraviolet to extreme ultraviolet. It took more than 10 years to develop these EUV lithography systems. It has been in production since 2019, first in TSMC and Samsung 7 nm nodes fabs, then at 5 nm from 2020. One key benefit of EUV etching is to reduce the usage of the costly multiple patterning process to improve lithography resolution. ASML's latest EUV lithography generation is dubbed High-NA (for high numerical aperture, 0.55 instead of 0.33) with an 8 nm resolution, using 13.5 nm EUV light and DSA (directed self-assembly)-based multi-patterning techniques to reduce the roughness of transistor patterns. That 13.5

some CMOS density technical challenges




Figure 10: Some of the key CMOS density technical challenges to overcome by the semiconductor industry. Sources: [26, 27].

nm light is not emitted by a laser directly. It comes from a plasma created by a high-power CO₂ laser vaporizing about 50,000 molten tin droplets per second.

A bit like in photography, High-NA optics will convey more light onto masks and silicon targets and will be required for nodes under 3 nm. It requires both new UV optics, and new light sources. And the EUV machines are much bigger and costly. These machines were deployed in 2024. The generation after High-NA would be Hyper-NA but even ASML is doubting it will be economically viable [28]. A current ASML EXE:5200B high-NA scanner costs about \$380M.

In December 2025, Reuters revealed that a state-backed consortium coordinated by Huawei had completed, in early 2025, a prototype EUV system in a high-security facility in Shenzhen. It is a system generating EUV light (13.5 nm) by vaporizing droplets of molten tin to create a light-emitting plasma, like with ASML, although the Chinese teams use solid-state lasers instead of the CO₂ lasers sourced by ASML from its Cymer subsidiary. Generating light is only the first part of a much more complicated system which must incorporate high-precision mirrors currently built only by Carl Zeiss (Germany) and are subject to export controls. Then, precision mechanics must move the wafers quickly with a nanometer-scale precision. The prototype fills an entire factory floor and has not exposed a single working wafer so far, with a chip production target set at 2028 at the earliest. So, creating a full-fledged EUV lithography production system in China will take several more years at best. A different Chinese effort is often confused with this one. Shanghai Aishengna Electronic Technology Group, a state-owned company created in August 2023 and identified by Reuters in July 2026, is mass-producing immersion DUV scanners and not EUV systems, with five machines planned for 2026 and twenty for 2027, to be compared with the roughly 130 immersion systems ASML ships every year.

Another avenue to improve lithography resolution is to use X rays instead of EUV. One US startup, Substrate, has ambitions in there, with a lot of technical and operational challenges ahead [29].

Reticles size corresponds to the optical systems used in lithography whose size is physically limited, especially optically. This limit has been reached with the largest recent processors.

Other scaling solutions were found including using vertical transistors like the traditional FinFET technology that has been in use for more than 10 years, that is now expanded with nanowires and nanosheets techniques as shown in Figure 11 [30], multi-die packaging associating multiple chips in a single packaging. The FD-SOI technology from CEA-Leti and STMicroelectronics adds an isolated layer of silicon oxide on silicon wafers, that limits the effects of transistor leakage and enables better operations at high frequencies with energy savings. It is particularly used in radio-frequency front-end chips in smartphones.

Transistor density fake news. After 2006, transistor density continued to grow. You've heard about these successive generations of 28 nm, 14 nm, 10 nm, 7 nm, 5 nm, 3 nm and now 2 nm transistor sizes. In May 2021, IBM announced it had prototyped 2 nm nanosheet-based chips, fabricated on 300 mm wafers at its own Albany NanoTech research complex and not in a Samsung fab, using EUV lithography [31]. IBM does not manufacture chips at volume anymore and licenses this technology to partners like Samsung and Rapidus. In December 2022, the company announced they could scale as low as 1 nm thanks to using ruthenium for chip interconnects [32]. In July 2021, Intel announced a new density scale using angstrom sized transistors, with 20 Å and 18 Å by 2025 (meaning... about 2 nm, given 1 Å = 0.1 nm). Intel eventually cancelled its 20A node in September 2024 to concentrate its resources on 18A. TSMC unveiled its N2 node roadmap in 2022 but its 2 nm volume production only started in Q4 2025, in its Hsinchu and Kaohsiung fabs, with a claimed 10% to 15% speed gain and a 25% to 30% power reduction over N3E, and a density increase of over 15%.

Unfortunately, this is misleading, probably one of the most significant “fake news” in the digital industry, and it has been going on for over 10 years. Seriously! These tiny transistors have no features with these announced




Figure 11: The various CMOS transistor technologies used as density increased. Nanosheets are becoming commonplace for transistor densities below the 5 nm marketing denomination. Source: [33].

sizes. This is a marketing trick from the whole semiconductor industry.

This is shown in Figure 12 with a table consolidated by the IEEE as of late 2023 as well as in IMEC’s roadmap in Figure 13. It describes all the transistor feature sizes for the “nodes” labelled 3 nm (2023) down to 0.5 nm or 5 angstroms that are planned for 2037. As mentioned in Wikichip technology node, “Since around 2017 node names have been entirely overtaken by marketing with some leading-edge foundries using node names ambiguously to represent slightly modified processes. Additionally, the size, density, and performance of the transistors among foundries no longer matches between foundries. For example, Intel’s 10 nm is comparable to foundries 7 nm while Intel’s 7 nm is comparable to foundries 5 nm”. What you discover here is that the metal pitch between “3 nm” transistors is of 24 nm and will reach 14 nm for “0.5 nm” transistors vs the 16 nm that were planned in IEEE’s 2022 roadmap. In these generations of transistors, the smallest features are the nanosheet thickness, which is 4 nm and the gate oxide thickness which is below 1 nm.

How is the semiconductor industry justifying this nodes marketing labeling? One is the real labeling is too complicated, with G48M24 for gate pitch and metal pitch sizes for “3 nm” densities. The second is that this fake density corresponds to the density power increase of these chips. Lastly, starting in 2028, “ground rule” transistor scaling is expected to slow down and saturate. Transistor density progress will then come from stacking several layers of transistors on top of the other, and improving backside metallization schemes, starting around 2031-2034. It will probably not avoid the heat barrier and the dark silicon phenomenon. In the end, the only feature that is below the 1 nm threshold is the gate oxide thickness, but it is a vertical, not a horizontal feature.

Other misunderstandings come from the confusion in density between the transistor features height and width. In April 2025, researchers in China published a Nature paper describing a new 2D RISC-V micro-controller chip using molybdenum disulfide transistors made with a sheet of molybdenum atoms sandwiched between two layers of sulfur atoms [35]. It was said that it could extend Moore’s law, without really discussing about

horizontal density [36]. But it can at least maybe have some positive impact on power consumption. Let alone the cost and time to industrialize any such new manufacturing technology. Other semiconductor innovations are geared toward saving energy, like this new bismuth based transistor also invented in China that could increase computing speed by 40% and save 10% energy. The difference however doesn’t seem significant enough to drive a whole semiconductor revolution [37].

1.5.4 Classical computing technology developments

The semiconductor industry used some other techniques to increase classical computing power, and all of them are not mentioned here. Since 2022, there’s been a big shift in the industry with a focus on supporting the training and inference of exponentially growing AI models. This means an increase of computing power but also speed in memory bandwidth. Another big concern for the semiconductor and classical hardware industry is to improve the power efficiency of these AI-supporting systems.

Domain Specific Architectures consist in encoding in various silicon features to make it more efficient both in speed and energy consumption. Most smartphone and laptop chips have been using and improving this technique for a while, embedding features like GPU cores, tensor cores for machine learning computing, audio and video codec DSPs, security units, input/output units and other spare features.

Multiple features are integrated in single die chips *aka* “system on chip” (SoC). It is even possible to create LLM dedicated DSA ASICs, like what Etched.ai (2022, USA, \$1.1B) is doing with its Sohu chip Figure 14, an inference-only transformer ASIC etched by TSMC in 4 nm, with over \$1B in signed contracts but no independent benchmark published so far.

2.5 and 3D packaging is another path used to integrate multiple features in small packaging associating specialized chips manufactured with different techniques (CPU, GPU, fast SRAM cache memory, storage, RF, photonic links features) and connected through high-speed links and buses. Nowadays large “chips” are actually chiplets made with several specialized chips,




Figure 12: The real transistor feature sizes per generation showing that 3 nm, 2 nm and below do not correspond to any real size in transistor designs in horizontal features. Starting in 2029, the generation naming will be in Angstrom, which is one tenth of a nanometer. Transistor size is not significantly changing from one generation to the other, validating the “end of Moore’s law” claim. The right way to describe these nodes would be a number scheme like G42M14/T6 in 2039 with a gate pitch of 42 nm, a metal pitch of 14 nm. It is of course more complicated than 1.8A! There are still 6 nm features like the spacers around transistor gates, and nanosheet thickness, which will shrink to 4 nm by 2031. Source: [34]. © Olivier Ezratty, September 2026, for the update and annotations.




Figure 13: IMEC roadmap in 2025 showing the evolution of transistor gate-pitch down to 10 to 14 nm in 2039. Another key technology progress comes from the integration of chip interconnect architecture which gets denser over time with back-side metal pitches decreasing from 65-160 nm down to 28 nm in 2025. Source: IMEC.

like with the Intel Ponte Vecchio and the Arrow Lake made with a GPU and SoC manufactured by TSMC and a CPU made by Intel. One example is the Intel Ponte Vecchio processor used in the DoE Argonne National Laboratory Aurora supercomputer. It is a chiplet containing 47 active chips [38].




Figure 14: The LLM-dedicated Etched board with its Sohu chip is packaged in rack servers for high-performance AI compute. Its HBM/SRAM hybrid memory design enabled faster compute power. Source: Etched.ai, 2026.

3D chips. In May 2026, at the IEEE ISCAS conference in Shanghai, Huawei co-president He Tingbo introduced the company's τ (Tau) Scaling Law which replaces geometric scaling with signal propagation time (τ) scaling

as the guiding principle for the evolution of semiconductors and electronic systems. It is documented in [39]. It favors architecture developments on speed vs transistor integration. The underlying approach named LogicFolding is a design methodology that partitions digital, analog, and memory circuits across 3D stacks, distributing critical-path gates across two and eventually more tiers connected by ultra-fine-pitch hybrid bonding. It focuses on optimizing the cost of data movement which consumes 80% of the energy in large AI clusters. Huawei claims it will reach a transistor density equivalent to a 1.4 nm node by 2031 while remaining stuck at around 7 nm in lithography and without any EUV tool. The technology is to be used in the new generations of the Kirin SoC for their smartphones, launched in the fall of 2026, and in the Ascend AI accelerators. Huawei says it has already mass-produced 381 chips based on these principles over the past six years. The density claim deserves the usual grain of salt since it relies on a density formula using a coefficient that inflates the result by about 36% compared with the industry-standard one, which would bring the Kirin 2026 back from 238 to about 175 million transistors per mm^2 , and since 3D EDA tooling for such stacks is not mature yet. But Huawei is not the only one working on 3D chips.

See for example, the prototype 3D chips created by Stanford, Carnegie Mellon, the University of Pennsylvania, MIT, and Skywater, in 2025 which includes silicon CMOS compute, RAM layer, carbon nanotube field-effect transistors on the same die, but in 90 nm to 130 nm process [40, 41].

Memory. One key technological development is to make sure that memory is as close as possible to processing units, in-memory processing being the extreme limit of memory and processing integration [23, 42, 43]. Improvements in cache management and other designs can potentially help gain an order of magnitude in computing speed, like with the technology from Blueshift Memory (UK) that currently targets RISC-V chips. Some memory progress happen which do not directly impact computing speed, like this non-volatile memory prototyped in China in 2025 which support read-write operations at 3 GB/s [44]. The fastest current memory is the SRAM that sits within the chips. It supports the cache memory. Then, you have external DRAM modules. The most common technology is HBM (High Bandwidth Memory), used by vendors like Nvidia in their GPU packaging. It is a 3D-stacked memory architecture designed to bypass the physical bandwidth limitations of traditional planar memory. The current production standard as of mid-2026, is HBM3e, which provides roughly 1.2 TB/s of bandwidth per memory stack using a 1024-bit interface. The new HBM4 standard doubles the interface width to 2048 bits, supporting 2.0 TB/s bandwidths. The subsequent HBM4e will reach 4.0 TB/s and HBM5 will double the interface to 4096 bits. Planned for 2032, HBM6 will reach 8.0 TB/s. Future GPUs will support eight HBM6 stacks and 64 TB/s of total memory bandwidth. Along a capacity of 1 TB, it will support entire LLM model or significant size.

HBM will soon also be completed by HBF memory (High Bandwidth Flash), proposed by SanDisk in 2025. HBF will sit next to HBM on future AI inference chiplets and complement the memory hierarchy [45–47]. HBM is for read-write tasks while HBF is adapted to read-only tasks which are heavily used in AI inferences.

There are different optimization approaches for handling the memory close to processors. Cerebras (USA) uses a zero-distance approach with memory physically distributed directly along the compute cores across an entire uncut silicon wafer (with 40 GB as of 2026). This larger internal SDRAM chip memory speed up communication within it but with the outside world. It uses slower DRAM access which is not built on HBMx standards like with Nvidia chips. Euclid (the Netherlands) uses dataflow-SIMD which operates directly on in-memory data streams with 8 PB/s of bandwidth within a highly condensed system-in-package (SIP). Vsora (France) with Jotunn and Tyr processor families uses a pipelined data flow with memory remaining off-chip, but with highly optimized internal pathways to ensure continuous data feeding. Vertical Compute (Belgium) stacks high-density MRAM (developed at Imec) on top of processing cores in a 3D chiplet architecture [48], reducing data travel distances from centimeters to nanometers between memory and compute, reducing its latency. Nvidia uses a more regular scaled Von Neumann architecture relying on banks of HBM memory packaged next to the GPU via 2.5D interposers.

SIMD, for single instruction multiple data processing, is used in GPUs and GPGPUs (general purpose GPUs). It is handling matrix multiplications in parallel. This is the technique used by Nvidia among others, starting with the “Volta” V100 in 2017, the “Ampere” A100 in 2020, the “Hopper” H100 in 2022, H200 in 2023, and the “Blackwell” B200 in 2024 [49]. The A100 had 54.2 billion transistors superseded closely in size by the Graphcore GC200 with its 59.4 billion transistors and 1,472 cores. The H100 launched in 2022 has 80 billion transistors, consolidating two adjacent chips in a single package.

Then came the Nvidia GH200 which embeds CPU arm cores, removing the need for a traditional CPU-to-GPU PCIe connection. This GPU used Nvidia NVLink-C2C chip interconnects, increasing the bandwidth between GPU and CPU by 7x compared with the latest PCIe technology and reducing interconnect power consumption by more than 5x. The H200 launched in November 2023 has 141 GB of HBM3E memory with a 4.8TB/s bandwidth.



Figure 17: A Nvidia NVL72 rack with its 72 GPUs and 36 CPUs. Source: Nvidia.

Unveiled at CES in January 2026 for deployments in the second half of the year, the Nvidia Rubin GPU chip hosts 336 billion transistors on two reticle-sized dies etched in TSMC 3 nm, which is 1.6 times Blackwell’s 208 billion. It uses HBM4 memory with 288 GB per package and a memory bandwidth of 22 TB/s, versus 8 TB/s for Blackwell, and delivers up to 50 PFLOPS of NVFP4 inference and 35 PFLOPS of NVFP4 training. Its TDP reaches about 2,300 W. The companion Vera CPU contains 88 Arm Olympus cores and 227 billion transistors, and is connected to the GPUs with NVLink-C2C at 1.8 TB/s. A standard Vera Rubin NVL72 data-center rack contains 72 Rubin GPU packages and 36 Vera CPUs, with 20.7 TB of HBM4 memory and a total of 75 TB of unified memory. It consumes about 220 kW (Figure 17). It uses the NVLink 6 Switch with up to 3.6 TB/s of GPU-to-GPU scale-up bandwidth per GPU, twice NVLink 5, and supports 3,600 PFLOPs FP4 inference capacity. The NVL72 is provided as a




Figure 15: How HBM memory chips are stacked on top of each other and connected to a chiplet through TSV (through-silicon vias), which itself connects to the GPU chip. Source: [47].




Figure 16: The future role of HBF memory (High Bandwidth Flash) which will complement HBM memory in the memory hierarchy, and enable faster AI inferences. Source: [47].

product by Nvidia (DGX) and a specification (MGX), supported by its OEM partners (Gigabyte, Asus, HPE, Supermicro, etc.)

TSMC plans to create 200 billion transistors monolithic chips and 1 trillion multi-die chiplets by 2030 [50, 51].

SSD storage with PCIe connectivity has accelerated computing by two orders of magnitude compared to classical hard disks. In your laptop, you can reach 7 GB/s with a PCIe 4.0 drive and up to 14 GB/s with PCIe 5.0, compared to about 100 MB/s with a hard drive. The integration levels in 3D NAND flash chips are like CMOS transistors with pitches that can go down to 12 nm. But since all memory is not used simultaneously, these chips used stacks of transistors. As of 2026, SK hynix mass-produces 321-layer chips, Kioxia and SanDisk sample a 332-layer BiCS10 generation, Samsung is ramping a 400-layer class V-NAND and has demonstrated a 900-layer research device. The 1,000 layers mark is in sight around 2030-2031 [52].

Chip size record can reach 21.5 cm × 21.5 cm. It was achieved in 2019 by Cerebras (USA), fitting its WSE-1 chip in an entire 300 mm wafer, which circumvents the reticle size limit by being etched in several runs, for its

84 identical dies or reticle fields, connected across the scribe lines by metal layers.

In March 2024, Cerebras introduced its new WSE-3 which contains 4 trillion transistors, 900,000 cores, 44 GB on-chip SRAM and is manufactured using a 5 nm TSMC process and run at 1.4 GHz. The new WSE-3 Turbo launched in August 2026 runs at double clock speed of 2.8 GHz and delivers 250 petaflops of sparse FP16 AI performance, or 125 petaflops dense. It can accommodate an external memory of 1.5 TB, 12 TB, or 1.2 PB. It can train LLM models with up to 24 trillion parameters, beyond what most LLMs have as of 2026. Manufacturing techniques generate defects and more than a couple percent of the 900,000 processing units are defective and are short-circuited during software execution. The massive Cerebras chip, shown in Figure 19, has a TDP of about 15 kW in its WSE-3 version, the Turbo pushing about twice this power through the same wafer to sustain its doubled clock. The dissipated heat is evacuated using a specific water-cooling system in their server.

The new full-rack CS-4 server was also launched in August 2026. It contains three WSE-3 Turbo chips and delivers 750 PFlops of AI compute, with 129.6 PB/s




Figure 18: Various 2D to 3D chips integration techniques into chiplets. The most recent 3D integration involved stacking different chips on top of each other and interconnect them with TSV (through silicon vias) vertical wiring within each chip. This technology is now being used in quantum computing, like with superconducting qubit chips. Source: [53].




Figure 19: The impressive Cerebras wafer-scale chip WSE-3 Turbo with its 4 trillion transistors and 44 GB on-chip cache memory, enabling very fast deep learning tasks thanks to its 2.8 GHz clock. Announced in August 2026, the CS-4 server rack accommodates 3 such chips. Source: Cerebras. Updated in September 2026.

memory bandwidth (43.2 PB/s per wafer), 160.5 PB/s fabric bandwidth, 7.2 Tb/s of system I/O bandwidth and a wafer-to-wafer interconnect latency down to 2 microseconds. It supports LLM inferences with 1,500 tokens per second on models like GPT 5.6 SOL.

Supercomputers can be built with thousands of such chips. In 2025, Cerebras deployed a supercomputer using 350 of their WSE-3 chips at Scale Datacenters, an AI datacenter company, in Oklahoma City (Figure 20).

On the use cases side, Cerebras and its partners and customers published interesting work on training large language models [55, 56], genetic algorithms [57], extending the time scale of simulable molecular dynamics [58], solving a 2D Ising model [59] and fluid mechanics simulations (with TotalEnergies [60]).

It adds an impressive 40 exaflops capacity of AI compute, consumes 10 MW and seems able to run large-scale AI inferences 20 to 50 times faster than equivalent GPU-



Figure 20: The Cerebras-based datacenter built in Oklahoma City in 2025 with 350 Cerebras chips and a peak performance of 40 exaflops of AI compute for a 10 MW power draw. Source: [54].

based datacenters [54]. It is no surprise the company could raise \$1.1B in a Series G in September 2025 at an \$8.1B post-money valuation, then \$1B in a Series H at about \$23B in early 2026, for a total pre-IPO funding of about \$2.8B, way more than PsiQuantum. Cerebras went public on the NASDAQ under the CBRS ticker on May 14, 2026, raising about \$5.5B at an implied valuation of around \$26B. Its revenue was \$79M in 2023, \$290M in 2024 and then \$510M in 2025, with a first net profit of \$88M, driven by a \$10B multi-year compute contract with OpenAI covering 750 MW of capacity.

Inline with Cerebras' wafer scale processor, Euclid's CRAFTWERK chip presented in September 2025 is about four times smaller in linear size, with an interposer of roughly 100×100 mm and "only" 16,384 SIMD units. Its computing power reaches 8 PFLOPS in FP16 or 32 PFLOPS in FP4, with a TDP of about 3 kW. It has 1 TB of custom ultra-bandwidth memory delivering an enormous 8,000 TB/s bandwidth. Configured in the rack mount Station CWS 32 using 32 of these chips consuming 125 kW, it would deliver 1.024 exaflops of AI compute, 32 TB of UBM and 7.68 million LLM token inferences per second. All these figures are modeled on Llama 4 Maverick for an architecture that is still in advanced design and has not been taped out, so they belong to the promises category for now.

With its D1 chipset presented in July 2021, Tesla chose another approach. Engraved in 7 nm, it has a computing capacity of 22.6 TFLOPS FP32, with 50 billion transistors and a 400W TDP. It contains 354 computing units with 1.25 MB SRAM per unit. They assemble these D1 in 25-chipsets tiles, consuming 15 kW, exactly like a Cerebras chipset.

1.5.5 Unconventional computing

In a dedicated part of this book, we will evaluate some the other avenues considered to overcome the current limitations of classical computing, which may provide some power or efficiency gains positioned between classical and quantum computing. These belong to the broad category of "unconventional computing".

This includes **superconducting computing** operating at low temperatures (investigated in the USA and Japan), **digital annealing** computing (proposed by **Fujitsu**), **neuromorphic computing**, **reversible** and/or **adiabatic computing** that could reduce energy consumption and circumvents Dennard's scale end, **probabilistic computing** as well as different breeds of **optical computing**.

The book will also delve into some of the inner workings of supercomputers and specialized processors to better understand their strengths and weaknesses. When comparing quantum computers to classical computers, we are better off knowing both sides of the equation, not just the loud new kid in town!

These are sort of backup solutions, should science fail to create scalable quantum computers. It will also complement quantum computing used in the context of hybrid computing. Interestingly, some unconventional computing avenues, such as superconducting electronics, are potentially enabling technologies for scaling certain types of quantum computers.

However, at this point, none of these solutions seem positioned to solve intractable problems although some of these are claiming they have this capacity, which is quite hard to fact-check at large scales.

The history of technology is about exploring multiple branches. Some fail. Some help each other. Also, some can suddenly wake up after being frozen for decades. The game is open!


Common quantum technologies misconceptions

Misconception	Explanation
Quantum computing is fast or instantaneous.	While quantum algorithms may be faster than their classical equivalents in a quantum advantage regime, they can be quite slow and run for hours, days if not months and more.
Quantum computing speedups are explained by the Hilbert space dimension of a qubit register (2^N).	The number of computational-basis states indeed scales exponentially with the number of qubits but it doesn't explain the potential quantum algorithms speedups. It depends on many factors like the type of qubit gates used, how the data is loaded in the quantum circuit, its degree of entanglement and, above all, the structure of the problem.
Quantum computers will soon break RSA keys and all Internet cybersecurity.	We'd need at least 100,000 physical qubits to break an RSA-2048 key in less than a year according to the latest 2026 estimates. We are many years away if not largely more than a decade away from this mark at best. So far, progress has been better with reducing quantum circuit size than improving hardware scaling.
Quantum computing will enable larger big data applications.	Not really. Quantum computers are slower by several orders of magnitude to feed with data compared to classical computers. This could slightly improve with the use of quantum sensor data or with advanced compression techniques.
Quantum computing will solve the LLM energy crisis.	Probably not. Loading both training data and a trained model in a quantum circuit would be prohibitively slow even for modest-size LLMs, and just for getting one response token!
Quantum computing will progress following some Moore's exponential law.	While qubit numbers are making progress and did fit some exponential trend for some time, and then stalled, it is not the case with qubit fidelities and many other figures of merit, particularly at scale. Researchers are still struggling to scale the number of qubits and keep good fidelities.
Quantum computing is business-ready now.	Many analysts and vendors would like you to believe it with NISQ, but it is not yet the case for industry grade applications. The few applications nearing quantum advantage are most of the time physics simulations far from industry use cases.
Quantum computing will (soon) help fix climate change.	It may enable some research in innovative chemical engineering, but this will require large-scale fault-tolerant quantum computers which are decades away. We'd better fix climate change with classical methods in the meantime.
China is investing between \$15B and \$25B in quantum technologies and is beating all other geographies.	This is probably not true or at least not verified nor decomposed in detail. Estimates range from \$5B to \$10B over 10 years. The \$15B to \$25B investment mark comes from a questionable 2017 statement on the Hefei lab buildup. This number has not changed since 2017! But China is faring well in many domains like in photonic quantum computing, quantum communications and quantum sensing.
Government funding and private investment means it will work.	Money can buy time, but the laws of physics make this hard. Creating viable quantum computers requires more iterations in component design and manufacturing, along a complex value chain.
Quantum communication will replace the Internet with faster transmission speed.	Quantum communication's main purpose is to enable the creation of symmetric encryption keys between trusted parties, not to speed up data communications. It will also enable the connection between multiple quantum computers by entanglement-based quantum communication links, but not faster than light. All these techniques need some classical links on top of quantum links.
Quantum computers will save energy.	We are not so sure about it, particularly with large scale fault-tolerant systems. This is an area of investigation and optimization.

Table 1: 12 misconceptions on quantum computing and communications © Olivier Ezratty, 2024-2026.

1.6 Why... key takeaways

- This book is unique in its shape, structure and length. It covers quantum technologies with a 360° approach. It is more scientific than most broad-reach publications, outside research review papers. It is a good appetizer for those who want to investigate the matter whatever the angle. It covers the broad field of quantum science and technologies with fundamental science and engineering viewpoints, but also through an economic and geopolitical lens. It contains an extensive bibliography with over 11,700 references, mostly scientific papers. It tries to answer many commonplace questions that are not well addressed in broad audience scientific and business publications.
- Many enabling devices used in existing digital technologies, including transistors and lasers, rely explicitly on quantum physics and are commonly associated with the “first quantum revolution”. They exploit light-matter interactions in large ensembles of quantum objects such as electrons, atoms and photons and in semiconductors. The “second quantum revolution” is an expression coined in 2003 for a newer generation of quantum technologies that relies on a variable mix of superposition, entanglement and control of individual quantum objects. It usually includes quantum computing, quantum telecommunications, quantum cryptography and quantum sensing. Quantum matter applications are a newer addition to this list that is also described in this book.
- Quantum technologies are at the crossroads of many scientific domains encompassing physics, mathematics, computing, social sciences and economics. This creates new educational and pedagogical challenges that must be addressed in innovative ways and customized for various audiences. This book targets broad audiences with a technical background, including computer science engineers, but also quantum physicists and quantum information scientists who want to have a broad view of the field and its burgeoning ecosystem.
- Quantum computing may accelerate selected problems for which suitable quantum algorithms and favorable end-to-end resource scaling are known. It does not make all classically intractable problems tractable. Enormous challenges must still be overcome to scale quantum computers beyond today’s relatively modest systems. In the interim, noisy intermediate-scale quantum computers (NISQ) may provide useful improvements for some specific tasks, while analog quantum computers are other candidates for demonstrating quantum advantage on selected problems.
- Moore’s law has not ended yet, particularly with regard to the number of transistors per chip. Classical computing continues to advance and remains much better suited than existing quantum computers to workloads involving large volumes of classical data. Data loading speed and big-data applications will remain key strengths of classical computing for the foreseeable future.
- Other new technologies may compete with quantum computing and belong to the broad “unconventional computing” category. Some may provide advantages for specific problem classes, including combinatorial optimization, but no generic exponential computing capacity should be inferred without a problem-specific complexity analysis and comparison with the best classical methods. Other approaches may instead provide benefits such as lower energy consumption. Technologies based on superconducting electronics and adiabatic or reversible computing could also help enable scalable quantum computing, particularly with superconducting qubits.
- At last, this part contains an inventory of 12 key misconceptions on quantum computing.



1.7 Bibliography and notes

These are the bibliographical references and notes for the **Why?** section.

- [1] [Le quantique, c'est fantastique](#), by Olivier Ezratty and Fanny Bouton, Web2day, YouTube, June 2018 (1h03mn).
- [2] [Conversations with the author in 2018 and onward](#).
- [3] [A Short Guide to Quantum Mechanics – Some Basic Principles](#), by Joachim Stolze, arXiv, August 2024 (155 pages).
- [4] [What Makes Quantum Computing So Hard to Explain?](#), by Scott Aaronson, June 2021.
- [5] [The Transistor, A Semi-Conductor Triode](#), by John Bardeen and Walter H. Brattain, Physical Review, July 1948 (2 pages).
- [6] [Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy](#), by John S. Bell, Cambridge University Press, 2nd edition, with an introduction by Alain Aspect, June 2004 (289 pages).
- [7] [Quantum technology: the second quantum revolution](#), by Jonathan P. Dowling and Gerard J. Milburn, Philosophical Transactions of the Royal Society A, August 2003 (20 pages)
- [8] [Quantum Technology Second Quantum Revolution](#), by Jonathan Dowling, 2011 (60 pages).
- [9] [The Second Quantum Revolution: From Entanglement to Quantum Computing and Other Super-Technologies](#), by Lars Jaeger, 2018 (331 pages).
- [10] [Diffraction Free Light Source for Ghost Imaging of Objects Viewed Through Obscuring Media](#), by Ronald Meyers, 2010 (22 pages).
- [11] [Quantum Imaging](#), by Yanhua Shih, arXiv, 2007 (25 pages).
- [12] [Opening the Black Box inside Grover's Algorithm](#), by E. Miles Stoudenmire and Xavier Waintal, PRX, November 2024.
- [13] [Scientists are using AI to dream up revolutionary new proteins](#), by Ewen Callaway, Nature, September 2022
- [14] [Value Migration: How to Think Several Moves Ahead of the Competition](#), by Adrian J. Slywotzky, 1995 (336 pages).
- [15] [Cramming more components onto integrated circuits](#), by Gordon Moore et al., Electronics, April 1965 (4 pages).
- [16] [Progress in Digital Integrated Electronics](#), by Gordon E. Moore, IEEE International Electron Devices Meeting, reprinted in IEEE Solid-State Circuits Society Newsletter, December 1975 (3 pages).
- [17] [Design of Ion-Implanted MOSFET's with Very Small Physical Dimensions](#), by Robert H. Dennard, Fritz H. Gaensslen, Hwa-Nien Yu, V. Leo Rideout, Ernest Bassous, and Andre R. LeBlanc, IEEE Journal of Solid-State Circuits, October 1974 (13 pages)
- [18] [Minimum Energy of Computing, Fundamental Considerations](#), by Victor Zhirnov, Ralph Cavin, and Luca Gamaitoni, 2014 (40 pages).
- [19] [World Records achieved with CPU Frequency](#), HWBOT.
- [20] [50 years of microprocessor technology trends](#), by Karl Rupp, February 2022.
- [21] [Implications of Historical Trends in the Electrical Efficiency of Computing](#), by Jonathan G. Koomey, Stephen Berard, Marla Sanchez, and Henry Wong, IEEE Annals of the History of Computing, July 2011 (9 pages)
- [22] [Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities](#), by Gene M. Amdahl, AFIPS Spring Joint Computer Conference, April 1967 (3 pages)
- [23] [Energy Efficient Computing Systems: Architectures, Abstractions and Modeling to Techniques and Standards](#), by Rajeev Muralidhar, Renata Borovica-Gajic, and Rajkumar Buyya, arXiv, March 2022 (63 pages).
- [24] [Ultimate physical limits to computation](#), by Seth Lloyd, arXiv, February 2000 (22 pages).
- [25] [Experimental verification of Landauer's principle linking information and thermodynamics](#), by Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz, Nature, March 2012 (3 pages) (preprint on [arXiv](#)).
- [26] [Reversible Circuits: Recent Accomplishments and Future Challenges for an Emerging Technology](#), by Rolf Drechsler et al., 2012 (8 pages).
- [27] [On the Vertically Stacked Gate-All-Around Nanosheet and Nanowire Transistor Scaling beyond the 5 nm Technology Node](#), by Hei Wong and Kuniyuki Kakushima, Nanomaterials, May 2022 (15 pages).
- [28] [Hyper-NA after high-NA? ASML CTO Van den Brink isn't convinced](#), by Saugata Ghose, Amirali Boroumand, Jeremie S. Kim, Juan Gómez-Luna, and Onur Mutlu, Bit Chips, September 2022.
- [29] [American startup Substrate promises 2nm-class chipmaking with particle accelerators, at a tenth of the cost of EUV — X-ray lithography system has potential to surpass ASML's EUV scanners](#), by Anton Shilov, Tom's Hardware, October 2025.
- [30] [Beyond CMOS, Superconductors, Spintronics, and More than Moore Enablers](#), by Jamil Kawa et al., March 2019 (43 slides).
- [31] [IBM Introduces the World's First 2-nm Node Chip](#), by Evan Ackerman, IEEE Spectrum, May 2021.
- [32] [The path to 1 nanometer chips and beyond](#), by Mike Murphy, December 2022.
- [33] [Beyond CMOS](#), by Jamil Kawa, March 2019 (43 slides).
- [34] [IEEE IRDS 2024: More Moore](#), by IEEE, 2023 (36 slides).
- [35] [A RISC-V 32-bit microprocessor based on two-dimensional semiconductors](#), by Mingrui Ao, Peng Zhou, et al., Nature, April 2025 (preprint on [Researching](#)).
- [36] [2D Chip Breakthrough: 6,000 Transistors, 3 Atoms Thick](#), by Charles Q. Choi, IEEE Spectrum, April 2025.
- [37] [Low-power 2D gate-all-around logics via epitaxial monolithic 3D integration](#), by Junchuan Tang, Hailin Peng, et al., Nature Materials, February 2025
- [38] [Intel's Take on the Next Wave of Moore's Law](#), by Samuel K. Moore, IEEE Spectrum, December 2022.
- [39] [A time scaling theory for multi-layer electronic systems](#), by Tingbo He, Science China Information Sciences, July 2026 (10 pages).
- [40] [Foundry Monolithic 3D Unlocks Large Throughput Benefits: 3D Memory with Tucked Sense Amplifiers + Logic using Heterogeneous Silicon CMOS + Resistive RAM + Carbon Nanotube FETs](#), by S. Choi et al., 2025 IEEE International Electron Devices Meeting (IEDM), December 2025
- [41] [First truly 3D chip fabbed at US foundry, features carbon nanotube transistors and RAM on a single die — future devices could have up to 1000x improvement in energy-delay product](#), by Luke James, Tom's Hardware, December 2025.
- [42] [Processing-in-memory: A workload-driven perspective](#), by S. Ghose et al., IBM of Research and Development, 2019 (19 pages).
- [43] [The Landscape of Compute-near-memory and Compute-in-memory: A Research and Commercial Overview](#), by Asif

- Ali Khan, Jeronimo Castrillon, et al., arXiv, January 2024 (42 pages).
- [44] [Subnanosecond flash memory enabled by 2D-enhanced hot-carrier injection](#), by Yutong Xiang, Peng Zhou, et al., Nature, April 2025.
 - [45] [H³: Hybrid Architecture Using High Bandwidth Memory and High Bandwidth Flash for Cost-Efficient LLM Inference](#), by Minh Ha, Euiseok Kim, and Hoshik Kim, IEEE Computer Architecture Letters, February 2026 (4 pages) .
 - [46] [Hot-Cold Tiering of HBM and High Bandwidth Flash for Agentic LLM Serving](#), by Jongjin Baek, Won Ji, Seungjae Yoo, and Joo-Young Kim, IEEE Computer Architecture Letters, August 2026 (4 pages)  (preprint on [arXiv](#)).
 - [47] [HBM-HBF Memory Hierarchy Innovations with Token Economics of Agentic AI Era](#), by Joungho Kim, KAIST TeraByte Interconnection and Package Laboratory, School of Electrical Engineering, June 2026 (130 slides).
 - [48] [Field-Free Spin–Orbit Torque Driven Switching of Perpendicular Magnetic Tunnel Junction through Bending Current](#), by Vaishnavi Kateel, Sebastien Couet, Kevin Garelo, et al., Nano Letters, June 2023 (8 pages).
 - [49] [Nvidia Unveils Blackwell, Its Next GPU](#), by Samuel K. Moore, IEEE Spectrum, March 2024.
 - [50] [How We'll Reach a 1 Trillion Transistor GPU](#), by Mark Liu, IEEE Spectrum, March 2024.
 - [51] [TSMC charts a course to trillion-transistor chips, eyes 1nm monolithic chips with 200 billion transistors](#), by Anton Shilov, December 2023.
 - [52] [Improvement of memory performance of 3-D NAND flash memory with retrograde channel doping](#), by Deepika Gupta, July 2023 (6 pages).
 - [53] [An introduction to chiplet-based architectures](#), by John Park, Chipscale Review, 2020 (4 pages).
 - [54] [Cerebras - new datacenter in Oklahoma City](#), by Andrew Feldman, Cerebras news, September 2025.
 - [55] [MediSwift: Efficient Sparse Pre-trained Biomedical Language Models](#), by Vithursan Thangarasa, Sean Lie, et al., arXiv, August 2024 (14 pages).
 - [56] [Who Has The Fastest AI Inference, And Why Does It Matter?](#), by Karl Freund, October 2024.
 - [57] [Trackable Agent-based Evolution Models at Wafer Scale](#), by Matthew Andres Moreno, Luis Zaman, et al., arXiv, June 2024 (14 pages).
 - [58] [Breaking the Molecular Dynamics Timescale Barrier Using a Wafer-Scale System](#), by Kylee Santos, Sivasankaran Rajamanickam, et al., arXiv, May 2024 (10 pages).
 - [59] [Record Acceleration of the Two-Dimensional Ising Model Using High-Performance Wafer Scale Engine](#), by Dirk Van Essendelft, Wissam A. Saidi, et al., arXiv, May 2024 (13 pages).
 - [60] [Massively Distributed Finite-Volume Flux Computation](#), by Ryuichi Sai, Randolph R. Settghost, et al., arXiv, April 2023 (10 pages).

This **Why?** section contains 60 bibliographical references and notes containing at least 1,581 pages.

2 History and scientists

After having set the stage, this section makes a historical detour to discover the origins of quantum physics. As with any scientific and technological endeavor, it is above all a great human story. Thus, the tribute paid here to the many scientists who, step by step, made all this possible, and the living ones who are still working on it. This part also explores the way the academic world operates for the newbie. It focuses on how scientific papers are being written, evaluated, published and communicated and how the academic world operates in general with its habits, conferences and the like.

Nanoscope physics. Quantum physics deals with atomic and sub-atomic level particles and with the interactions between electromagnetic waves and matter (Figure 21). It differs from classical Newtonian physics, which predictably governs the dynamics of macrophysical objects, beyond a few microns and up to the size of planets and stars. Classical physics is governed by Newton's laws for matter, by Maxwell's laws for electromagnetic fields and associated forces and by statistical physics which describes continuous media such as gases and fluids and from which the principles of thermodynamics are derived.

When the speed of objects becomes close to the speed of light or when we reach large object's mass, the theory of relativity comes in, explaining the curvature of space-time and modelling the impact of gravity. It helps describe extreme phenomena such as black holes or neutron stars. It allows us to interpret the History of the Universe, but not entirely. But relativistic electrons are also hidden in our body's atoms and in many elements on earth as we'll quickly discover with the weird field of relativistic quantum chemistry.

The fourth domain of physics in this quadrant is quantum fields theory. It describes the physics of high-speed elementary particles, such as those observed in particle accelerators like quarks and the famous Higgs boson. Richard Feynman is one of the founders of quantum electrodynamics, a subset of quantum field theory.

In a way, quantum physics was a means to unify classical matter physics and electromagnetic waves physics. It helps describe how matter was organized at the atomic and electrons levels and how these interacted with quantized electromagnetic waves, aka photons, including visible light.

Unification is still in the making. Physics is still not yet complete nor unified. Some observable physical phenomena still resist it. We do not know how to explain the origins of gravitation and we are still looking for the dark matter and energy that would explain the cohesion of galaxies and the Universe current expansion. Scientists would like to explain everything, but some knowledge may never be accessible such as the shape and form of the Universe before the Big Bang.

The so-called theory of everything (ToE) or unification theory sought after by some physicists would be

a formalism unifying all theories of physics and in particular relativity, gravity and quantum physics. This very serious field of physics is still in the making. The American-Japanese physicist Michio Kaku estimates that some theory of everything will be finalized by 2100 [61]. He is at the origin of string theory. He defines the connection between the different branches of physics and this theory of everything [62]. But he happens to be overly optimistic in his prediction, particularly given the inaccessible sheer size and scale of particle accelerators that would be needed to expand the existing field of high-energy particle physics. It creates a crisis with physicists in that field who wonder what to do next.

Numerous proposals emerge and sorting them out is not easy. This is the case of the Wolfram Physics Project launched in April 2020 by Stephen Wolfram, a prolific Anglo-American physicist, mathematician and computer scientist. Building on his 2002 book [63], the author's idea is to explain everything, the world, physics, the universe, whatever, with cellular automata, graphs and fractals [64, 65]. The world would be discrete on a small scale, including time. His Physics Project focuses on the unification of physics with the same set of tools. Physicists' views on this theory are more than circumspect. The paper does not develop a theory that would be verifiable or falsifiable with an experimental approach as was the case for quantum physics (superposition, wave function, wave function collapse, atomic transition spectral lines, ...). Wolfram's theory was critically analyzed by various scientists [66, 67].

The same "hammer/nail explains everything" approach was created by a team of scientists who describe the Universe physics laws self-learning capabilities with a giant neural network approach [68].

Connecting the dots. This part will help you memorize who's who in the History of quantum physics and quantum computing. It will also cover some important science basics such as the Maxwell and Schrödinger equations that I will try to explain in layman's terms, at least for readers having basic sciences knowledge. Explaining quantum computing inevitably starts with some quantum physics 101 explanations. Some of its basics, although sometimes quite abstract, must be understood. I still always try to connect the dots between quantum physics and quantum computing from a practical basis. It is a vast puzzle. I will add its pieces one




Figure 21: High-level classification of the branches of physics. © Olivier Ezratty, 2020.

by one and even though the puzzle may not be fully completed, you'll get a picture enabling you to become well educated in quantum computing.

Experiments and theories. Quantum physics took shape in 1900. Like almost all sciences, it is the result of the incremental work of many scientists with interactions between experimentation, theories and mathematical creativity. Sometimes, quantum physics is better explained with its underlying mathematical models than with incomplete physical interpretations. Representation models such as the broad field of linear algebra play a key role in describing quantum states and their evolution in space and time. Linear algebra is also an essential tool to understand how quantum computer qubits are manipulated and measured. Even if we can trace the beginning of quantum physics to Max Planck's 1900 quanta discovery, it was based on earlier work from many other scientists who devised about the particle or wave nature of light, on the discovery of electromagnetism and atoms. Quantum physics is a human adventure that brought together immense talents who confronted each other and evolved step by step their understanding of the nanoscopic world. New generations of scientists have always questioned the state of the art built by their predecessors. Max Planck's cynically explained in 1950 the evolution of science with the death of old generation of scientists: "*A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die and a new generation grows up that is familiar with it*".

Physicists conducted numerous experiments, elaborated theories, and then verified them experimentally, sometimes with several decades of latency. They also had to pour philosophy into their work to interpret the deep significance of their discoveries, and quantum physics was not an exception. Despite its constant enrichment, quantum physics has shown an astonishing robustness to stand the test of time and with extreme precision, in the 10^{-12} range.

Misrepresentations. Many quantum physics scientists are famous even for general audiences, even though their work has been overly simplified. Schrödinger's famous cat and Heisenberg's indeterminacy principle are commonplace... even when their underlying details are quite different from their related clichés. Schrödinger's key work is his non-relativistic particles wave equation, not the ten or so lines he wrote in 1935 on his eponymous cat thought experiment that is usually grossly misinterpreted!

Like life in general, science is a great relay race, with many players. Hundreds of other less-known contributors have also grown the field and must be recognized. Sometimes, genius scientists were so prolific that we forget their contributions. This is the case of John von Neumann who is better-known for his "von Neumann model" that is the cornerstone of classical computing and for his contribution to the design of EDVAC between 1945 and 1949, the first stored program computer architecture, rather than for his huge contribution to quantum physics mathematical formalism with density matrices and quantum measurement. It depends on the field you are working in, classical computing or quantum physics.

You will not find here inventors or entrepreneurs *ala* Steve Jobs or Elon Musk, even though the founders of startups like D-Wave, IonQ, Rigetti and PsiQuantum are among the entrepreneurial pioneers of this burgeoning industry, all being high-level scientists with a PhD!

Hall of fame. The History of 20th century quantum physics is embodied in the mythical Fifth **Solvay Conference** in 1927, held at the Institute of Physiology in Brussels (Figure 22). It brought together the greatest mathematicians and physicists of the time including almost all the historical founders of quantum physics with Max Planck, Albert Einstein, Niels Bohr, Louis de Broglie, Erwin Schrödinger, Max Born, Werner Heisenberg and Paul Dirac. It was the first Solvay conference organized after World War I that welcomed German physicists. All this happened as the foundations of 20th

quantum physics theories were well laid out. 17 of its 29 participants got a Nobel Prize, 6 of which before the congress (names underlined in green) and the others afterwards (in blue). It was probably one of the largest concentrations and density of scientific brains per square meter in the history of mankind. It was the stage of a heated debate on the interpretation of quantum physics which later led to the diverse field of “quantum interpretations” and “quantum foundations” [69, 70]!

Solvay conferences on physics have been held every 3 to 4 years since their creation in 1911 by the entrepreneur and chemist **Ernest Solvay**. The 1927 congress’s topic was electrons and photons, which are at the heart of quantum physics. Half of these conferences are dedicated to quantum physics, the other on different branches of physics. The 28th edition was held in May 2022 and gathered a contemporary hall of fame of quantum scientists from quantum physics to quantum information science.

Here, the major contributions of early scientists in quantum physics are generally arranged in chronological order, with some indication of who influenced who.

2.1 Precursors

We begin with the classical physicists and mathematicians of the 18th and 19th centuries who laid the scientific groundwork that allowed their 20th century successors to formalize the foundations of quantum physics (Figure 23). It is roughly organized in scientific contributions chronological order.



Thomas Young (1773-1829, English) was one of the greatest sciences and arts polymaths of his time, working in optics, medicine, linguistics, Egyptology, and music. He determined that light behaved like a wave, which he proved with the double-slit experiment around 1801.

It is illustrated in Figure 24 and now bears his name. When reducing the size of both slits, it generates interference fringes creating alternating light and dark zones related to the wave nature of light. We had to wait till Albert Einstein’s work in 1905 to determine that light was also made of particles. His experiment used red filtered sunlight going through a first slit. Contemporary experiments use coherent laser light sources. This experiment is one of the cornerstones leading much later to the creation of the electromagnetism theory by James Maxwell. The slit experiment was implemented with electrons in 1961, with a similar result, illustrating the electron wave-particle duality, devised first by Louis de Broglie in 1924. It was then also done with atoms in 1991 and with various molecules starting in 2002.

Thomas Young also worked on the principles of refraction and human trichromatic vision as well as in fluid mechanics, including on the notion of capillarity and

surface tension. As an Egyptologist, Thomas Young contributed to the study of the hieroglyphs of the famous Rosetta Stone, which was later used by **Jean-François Champollion** to decipher the whole stone texts. Champollion was then sponsored and helped by a certain **Joseph Fourier**. Yes, the mathematician and physicist!



William Rowan Hamilton (1805-1865, Irish) was a mathematician and astronomer. He invented around 1827 a set of new mathematical formulations of the laws of physics incorporating electromagnetism.

Hamiltonians or Hamiltonian functions are all over the place in quantum physics. These are mathematical operators used to evaluate the total energy of a system of elementary particles including their kinetic and potential energies. This energy is evaluated over time. The Hamiltonian was first introduced in 1811 by Joseph-Louis Lagrange as “the Legendre transformation of the Lagrangian”. Hamiltonian owes its name to the later development of Hamiltonian mechanics, a reformulation of Lagrangian classical mechanics that Hamilton laid out in his two memoirs “On a General Method in Dynamics” published in 1834 and 1835. Hamiltonian mechanics has a close relationship with geometry, particularly symplectic geometry and Poisson structures. It created a link between classical and quantum mechanics. The Hamiltonian provides a mathematical framework for describing the dynamics and evolution of quantum states. Schrödinger’s 1926 wave equation describes the evolution of a system’s Hamiltonian over time. This concept is also used in analog quantum computing with quantum simulators and quantum annealers, like with D-Wave’s systems.

Hamilton is also behind the creation of quaternions in 1843 which generalize complex numbers, using i , j and k as imaginary numbers with $i^2 = j^2 = k^2 = ijk = -1$. It can be used to compute three-dimensional rotations and have some applications in quantum computing like for the representation of two-qubit entanglement and of single qubit gates from the Pauli group, in topological quantum computing. This is an exotic domain that we won’t cover in this book.



Niels Henrik Abel (1802-1829, Norwegian) is a prolific mathematician at the origin of the so-called Abelian groups. He died at the age of 26 from tuberculosis at Froland in Norway while visiting his fiancée, two days before a letter arrived offering him a professorship in Berlin.

His work focused on the semi-convergence of numerical series, sequences and series of functions, the convergence criteria of generalized integrals, the notion of elliptic functions and integrals (used in cryptography) and the resolution of algebraic equations including his proof of the impossibility of solving general quintic equations. Along with William Rowan Hamilton, Charles Hermite



Figure 22: The famous Solvay 1927 conference photo with its 17 Nobel prize laureates (6 back then underlined in green, and 11 after the conference, underlined in blue). It was the stage of heated debates between Niels Bohr and Albert Einstein on the completeness or incompleteness of quantum mechanics. Photo credit: Benjamin Couprie, Institut International de Physique de Solvay.



Figure 23: Precursor scientists who laid the ground particularly in the electromagnetic fields (green) and mathematics domains (blue).

and Emmy Noether, he is one of the main ‘suppliers’ of the mathematical foundations used in quantum mechanics. The adjectives “Abelian” and “non-Abelian” are associated with anyons, the quasiparticles that are the basis of topological quantum computing.

Why do these concepts of quantum mechanics invented long after his death refer to this mathematician? Mainly because the distinction between Abelian and non-Abelian is linked to their commutative mathematical representation. A system with A and B is Abelian when $A*B = B*A$, or non-commutative and non-Abelian when $A*B$ is not equal to $B*A$. In quantum physics, non-commutativity involves square matrices with the same size, starting with the Pauli matrices for which XY is not equal to YX . Non-commutativity is frequently found in quantum physics and particularly with quantum measurement. The order in which quantum objects properties are measured may influence the

results because the used measurement operators are non-commutative. In some cases, though, operators are commutative, like with the Measurement-Based Quantum Computing (MBQC) technique that we will have the opportunity to describe later when dealing with photon-based quantum systems.

A cruder illustration of mathematical non-commutativity uses non-square matrices, where the two orders do not even produce matrices of the same shape. Indeed, the multiplication of a matrix $(p \times q) * (q \times p)$ will give a matrix $(p \times p)$ whereas in the other direction, $(q \times p) * (p \times q)$ will generate a matrix $(q \times q)$, q and p being here a number of rows and columns.




Figure 24: The double-slit experiment principle © Olivier Ezratty, various sources compilation



Charles Hermite (1822-1901, French) was another prolific 19th century mathematician. He worked on numbers theory, quadratic forms, the theory of invariants, orthogonal polynomials, elliptic functions and algebra.

His main works were concentrated on the 1848-1860 period. We owe him the notion of Hermitian functions and matrices, which are widely used in quantum physics and quantum computing. A Hermitian matrix is composed of real numbers in the diagonal and can be complex in the rest and is equal to its transconjugate.

Namely, their transpose matrix whose value of complex numbers has been inverted (i becomes -i and vice-versa) as shown in Figure 25.

Quantum measurement operations in quantum physics and computers are defined by Hermitian matrices.

$$A = \begin{bmatrix} 2 & i & -2i \\ -i & 1 & 3 \\ 2i & 3 & -1 \end{bmatrix} \quad A^T = \begin{bmatrix} 2 & -i & 2i \\ i & 1 & 3 \\ -2i & 3 & -1 \end{bmatrix}$$

transposed matrix

$$A = A^\dagger = \overline{A^T} = \bar{A}^T$$

Hermitian matrix
matrix equals its transconjugate

Figure 25: How a Hermitian matrix is constructed.

Achille Marie Gaston Floquet (1847-1920, French) was a mathematician who developed mathematical analysis in the theory of differential equations. His name is used in Floquet codes (quantum error correction codes) and we also find him in the physics of quantum matter.



James Clerk Maxwell (1831-1879, Scottish) created in 1865 the theory of electromagnetic fields, combining an electric field and a magnetic field orthogonal to the direction of wave propagation, and moving at the speed of light.

This theory explains light-light interactions such as reflection, diffraction, refraction and interferences. Maxwell's work built on and improved the formalism created by Michael Faraday, Carl Friedrich Gauss, and André-Marie Ampère.

An electromagnetic field is produced by electrically charged objects with two interrelated components: an electric field created by stationary or moving electric charges and a magnetic field generated by moving electric charges. Maxwell's equations illustrate that when they are constant, electric, and magnetic fields are independent, and in variable regime (with a wavelength λ), it becomes interdependent ($\vec{E}$ and $\vec{B}$), one generating the other and vice-versa, hence the notion of electromagnetic waves and fields (Figure 26).




Figure 26: Electromagnetic wave's electric and magnetic fields components.

In Maxwell's equations, the electromagnetic field is represented by an electromagnetic tensor, a 4×4 matrix whose diagonal is zero and whose half of the compo-

nents describe the electric field and the other half the magnetic field. These four dimensions correspond to space (3) and time (1).

We now use four main Maxwell equations:

- The **Maxwell-Gauss** equation describes how an electric field is generated by electric charges. At each point in space, the electric field is directed from positive to negative charges in directions depending on the charges space position (Figure 26).

$$\text{div}(\vec{E}) = \frac{\rho}{\epsilon_0}$$

Figure 27: Maxwell-Gauss equation describing the electric field created by electric charges.

- The **Maxwell-flux** equation states that a magnetic field is always generated by a dipole with positive and negative charges that are connected and inseparable. Mathematically, this translates into the fact that the divergence of the magnetic field is zero and that there is no magnetic monopole (Figure 28), which means that there is no magnetic field line that escapes to infinity as we have with an electric field.

$$\text{div}(\vec{B}) = 0$$

$$\oint_{(\Sigma)} \vec{B} \cdot d\vec{S} = 0$$

Figure 28: Maxwell-flux equation.

- The **Maxwell-Faraday** equation describes how the variation of a magnetic field creates an electric field. This is the principle used in electric alternators. The rotational operator using a nabla sign ∇ corresponds to a differential vector operation. It is equal to the first derivative of the magnetic field over time (Figure 29).

$$\vec{\nabla} \times \vec{E} = - \frac{\partial \vec{B}}{\partial t}$$

Figure 29: Maxwell-Faraday equation connecting the magnetic and electric fields.

- The **Maxwell-Ampere** equation states that magnetic fields are generated by electric currents or by

the variation of an electric field. This interdependence between magnetic fields and varying electric fields explains the circulation of self-sustaining electromagnetic waves. Figure 30 shows the equation is the rotational magnetic field.

$$\vec{\nabla} \times \vec{B} = \mu_0 \vec{J} + \mu_0 \epsilon_0 \frac{\partial \vec{E}}{\partial t}$$

Figure 30: Maxwell-Ampere equation connecting magnetic field to electric field.

As with Schrödinger's equation, Maxwell's equations have several variations, which may be confusing. Maxwell first published twenty equations with twenty unknown variables in 1865. In 1873, he reduced them to eight equations. In 1884, **Oliver Heaviside** (1850-1925, English) and **Willard Gibbs** (1839-1903, American) downsized the whole stuff to the four partial differential vector equations mentioned above. These four vector equations are reduced to two tensor equations for electromagnetic waves propagated in vacuum (Figure 31). The non-interaction with other elements explains the independence in this equation between electric and magnetic fields.

Maxwell predicted that electromagnetic waves were travelling at the speed of light.

Electromagnetic waves were only experimentally discovered after Maxwell's death, by **Heinrich Hertz** (1857-1894) between 1886 and 1888. Hertz also discovered the photoelectric effect in 1887. Maxwell's description of electromagnetic waves had a phenomenal impact on electromagnetic telecommunications and optronics. It also served as a foundation for the first quantum physics laws developed by Max Planck in 1900 which led progressively to the quantization of the electromagnetic waves.

$$\frac{1}{c_0^2} \frac{\partial^2 \vec{E}}{\partial t^2} - \nabla^2 \vec{E} = 0$$

$$\frac{1}{c_0^2} \frac{\partial^2 \vec{B}}{\partial t^2} - \nabla^2 \vec{B} = 0$$

Figure 31: Maxwell's equations in vacuum.

Maxwell is also at the origin of the **Maxwell-Boltzmann** statistical law of gas distribution. It models the particle velocity distribution of a perfect gas. It does not take into account the interactions between particles and is not applicable to extreme conditions, such as very low temperatures.

In particular, it is replaced by the **Bose-Einstein** statistics for bosons (integer spin particles such as helium 4, which can be gathered in the same quantum state and energy level) and by the **Fermi-Dirac** statistic for fermions (particles with half-integer spins such as electrons or helium-3, which cannot cohabit in the same quantum and energy state).




Figure 32: Maxwell's demon principle. Source: Wikipedia.

Maxwell is the designer in 1867 of the so-called **Maxwell's demon** thought experiment which would make possible the reversibility of thermodynamic exchange processes and invalidate the second law of thermodynamics [71]. In Maxwell's original setting, a single vessel filled with a gas at uniform temperature is divided into two chambers by a partition pierced with a hole and a closure controlled by a "demon". Fast and slow molecules are freely mixed on both sides. Once mixed (Figure 32), the demon would sort them, letting the faster molecules through in one direction and the slower ones in the other, taking advantage of the natural spread of the gas kinetic energy. This would allow in theory and after a certain time to recreate a temperature difference out of an equilibrium situation (on the right in Figure 32), apparently violating the second law.

It took several decades to find the fault, notably via Léo Szilard in 1929 and Léon Brillouin in 1951. Initially, the explanation was that the demon needs to consume some energy to obtain information about the state of the gas molecules to sort them out. Therefore, energy is consumed to modify the stable equilibrium obtained to mix the gases [72].

The "up to date" explanation is somewhat different. The energy cost comes from resetting the demon's memory, which ultimately consists of a single bit of information. Here is the detailed explanation: we can understand the operation of resetting a bit of memory by considering an ultimate Carnot engine, consisting of a single particle that can be located to the left or right of a certain thermostated volume. Left = 0, Right = 1 There are two possible operations. The first one is compression. The particle is initially to the left or to the right of the volume that contains it (we don't know) and we compress the said volume so that at the end it is necessarily on

the left. It is an initialization operation where the bit is reset to state 0. As with any compression, you have to pay, here in this ultimate case, the work to be expended which is $kT \log 2$. This is Landauer's famous work, which sets an energy bound to all logically irreversible operations. The second operation is relaxation. In the beginning, we know whether the particle is on the left or on the right. We position a wall, a pulley with a mass at the end and let the trigger operate while extracting an elementary work equivalent to $kT \log 2$. This is a Szilard engine.

These two manipulations were performed experimentally at ENS Lyon by Antoine Bérut, Sergio Ciliberto et al. and published in Nature in March 2012, showing the energy footprint of information, and they are the ultimate solution to Maxwell's demon paradox [25, 73].

Another explanation positions the cost of Maxwell's demon in the measurement apparatus [74]. In 2025, Maria Violaris from the University of Oxford also proposed a way to test a Szilard engine with qubits [75].

All this had repercussions on the notion of the energy value of information and led, much later, to the creation of the field of information thermodynamics, i.e., the study of the energetic and entropic footprints of information, particularly in quantum computing.

This field was then investigated by **Rolf Landauer**, known for his study of irreversible information management circuits heat generation, and by **Charles Bennett** and **Gilles Brassard**, the co-inventors of the QKD-based BB84 protocol, which we will discuss later, and then by **Paul Benioff**, who was at the origin of the idea of gate-based quantum computing.

We finally owe Maxwell the creation of color photography in 1855, that was implemented in 1861, based on the three primary colors of human vision.

Maxwell's electromagnetic field equations have very well survived the test of time. It is still the basis of classical optics and quantum optics. Even when studying quantized light, researchers and students still rely on Maxwell's equations and their subsequent derivations created since then.



Ludwig Boltzmann (1844-1906, Austrian) was a physicist, the father of statistical physics, defender of the existence of atoms, facing a strong opposition from scientists until the beginning of the 20th century.

He created equations describing fluid and gas dynamics in 1872. He is also at the origin of the probabilistic interpretation of the second law of thermodynamics, which establishes the irreversibility of physical phenomena, particularly during thermal exchanges.

Irreversibility is associated with the creation of entropy which measures the level of disorder in a system. Boltzmann's equation *aka* the Boltzmann-Planck equation relates the entropy S of an ideal gas to W , i.e., an in-

teger number of microstates corresponding to the gas's macrostate: $S = k_B \log(W)$ where k_B is the Boltzmann constant, 1.380649×10^{-23} J/K. A number of microstates correspond to the different possible arrangements of molecular position and kinetic energy at a particular thermodynamic state.

Boltzmann tried his hand at philosophy while defending the existence of atoms. Depressed, he died by committing suicide.



Henri Poincaré (1854-1912, French) was a mathematician and physicist, precursor of the theory of relativity and gravitational waves. He was a philosopher of sciences and one of the last « universal scientist » covering many branches in mathematics, physics and philosophy.

We owe him the Poincaré sphere, a geometric representation of the polarization states of light which is the optical equivalent of the Bloch sphere that we will see later, and which mathematically describes the state of qubits. He is also the author in 1904 of the mathematical conjecture that bears his name, proved by Grigori Perelman in a series of preprints posted in 2002 and 2003 and validated by the mathematical community in 2006. It is relative to the 3-sphere, the hypersphere that bounds the unit ball in a 4-dimensional space.

He was also a first cousin of Raymond Poincaré (1860-1934), president of France during the First World War, a lesser-known figure than Georges Clémenceau who was then Prime Minister and drove the war efforts against Germany and with allies from the UK and the USA.



David Hilbert (1862-1943, German) is yet another prolific mathematician who, at the end of the 19th century, was the creator of the mathematical foundations widely used in quantum physics.

It covers in particular the Hilbert spaces of abstract vectors and generalized notions of lengths, angles and orthogonality. Hilbert spaces are used to represent the state of quantum objects and qubits with vectors and complex numbers, with inner products, distances and orthonormal bases (see Figure 33). His work had nothing to do with quantum physics, which was not yet formulated at the time. It is now commonly used in quantum physics' mathematical formalism. This is a common phenomenon in mathematics. Another example is the creation of non-Euclidean geometry in the 19th century by Carl Friedrich Gauss, Nikolai Ivanovich Lobachevsky, and Georg Friedrich Bernhard Riemann which was later used in general relativity by Albert Einstein.

His work was used by Paul Dirac in 1930 and John von Neumann in 1932 to lay the groundworks of quantum physics mathematical foundations like the Dirac Bra-Ket notation and the von Neumann quantum measurement formalism.




Figure 33: A Hilbert space is a vector space with an inner product. It enables the measurement of vector distances, angles, and lengths. Source: compilation Olivier Ezratty, 2022.



Joseph John Thomson (1856–1940, British) was a physicist who studied cathode rays, which led him to discover in 1897 the existence of the electron, a subatomic particle with a negative electric charge.

He was awarded the Nobel prize in physics in 1906. He also found the first evidence for isotopes of a non-radioactive element, neon, which he reported in 1913. Many of these non-radioactive isotopes are commonplace in quantum technologies like ^3He (used in cryostats), ^{28}Si (in silicon qubits), ^{12}C (in carbon nanotube and color centers qubits) and ^{87}Rb (with cold atom qubits).



Pieter Zeeman (1865-1943, Dutch) was a physicist, Nobel prize in Physics in 1902 with Hendrik Lorentz, for the discovery of the effect that bears his name between 1896 and 1897. The Zeeman effect occurs when excited atoms are exposed to a magnetic field.

This affects their emission or absorption spectrum, that displays many discrete spectral lines. The effect is observed with spectroscopy, which breaks down light rays of different wavelengths with a prism.




Figure 34: Normal Zeeman's effect energy transitions. It comes from the interaction between an external magnetic field and the orbital magnetic moment of electrons in an atom. Source: [76].

In his experiment, spectral lines are broken down into an odd number of lines (normal Zeeman effect, as shown in Figure 34 for cadmium atoms) or an even number of lines (abnormal Zeeman effect). The decomposition depends on the intensity of the magnetic field passing through the analyzed atoms. There is also a nuclear Zeeman effect explained by the spin of atom nucleus. It is matched by a polarization of the generated light whose nature and intensity depends on the orientation of the magnetic field relative to the light beam.

The normal Zeeman effect is explained by transitions between electron states having different orbital angular momentum, while the anomalous effect requires the electron spin identified in 1925 by George Uhlenbeck and Samuel Goudsmit and the Landé g-factor which combines spin and orbital contributions. Pauli's exclusion principle plays no direct role here. In astrophysics, the Zeeman effect measurement is used to evaluate the intense magnetic fields in stars as well as within the Milky Way. The nuclear Zeeman effect is used in magnetic resonance spectroscopy in MRI scanners.



Hendrik Antoon Lorentz (1853-1928, Dutch) was a physicist who worked on the nature of light and the constitution of matter and made the link between light and Maxwell's electromagnetism equations.

We owe him the Lorentz transformations that explain the results of Michelson-Morley's experiments between 1881 and 1887 which showed that the speed of light is stable, whatever the reference frame. With Henri Poincaré and George Francis FitzGerald (1851-1901, Irish), he was a key contributor to the theory of relativity formalized later by Albert Einstein between 1905 and 1915.

Let's also add **Joseph Larmor** (1857-1942, Irish-British) who, among other various contributions, was one of the first to associate electric charges with electron particles in 1894. We also owe him the notion of Larmor precession, the rotation of the magnetic moment of an object when it is exposed to an external magnetic field, which he derived theoretically in 1897. Its experimental exploitation came much later with nuclear magnetic resonance on protons, from Isidor Rabi in 1938 to Felix Bloch and Edward Purcell in 1946.

2.2 Founders

The foundations of quantum physics started with Max Planck's black-body explanation with energy quanta and, then took shape over three and a half decades, roughly until 1935 (Figure 35). It involved the successive contributions from Einstein, Bohr, De Broglie, Schrödinger, Born, Heisenberg, Dirac, and von Neumann to mention only the best-known contributors who were all theoreticians and not experimentalists.

Things were relatively quiet during World War II as lots of scientists were focused on creating the atomic bomb in the USA under the umbrella of the then very secret Manhattan project while Europe was not the best place in the world for travel and international scientific collaborations.

German scientists who initially led quantum physics became isolated or emigrated to the USA or the UK because they were Jews, like Albert Einstein or Max Born.

Here is a broader tour of the great physicists and mathematicians who laid the foundations of quantum physics. They are all Europeans, many of whom emigrated from Europe to the United States before World War II.



Max Planck (1858-1947, German) was a physicist, initially specialized in thermodynamics. In 1900, he developed the first basis of quantum physics, hypothesizing that energy exchanges between light and matter are made by discrete quanta.

This radiation is not continuous but varies by thresholds, in steps of a certain amount of energy, hence the term "quantum" and "quantum physics" or "quantum mechanics". His theory allowed him to roughly explain for the first time the enigmatic radiation of black bodies, that absorbs all incident magnetic radiation (Figure 36). Examples of black bodies are a closed cavity like an oven, a heated metal that becomes red, orange, then white depending on the temperature, or a star like our own Sun. The spectrum of electromagnetic waves emitted by a black body depends only on its temperature and not at all on its material.

The higher the temperature is, the more the electromagnetic spectrum emitted by the black body slides towards higher frequencies on the left, therefore towards purple and ultraviolet. The theory solved the ultraviolet catastrophe.

This "ultraviolet catastrophe", an expression created later in 1911 by **Paul Ehrenfest** (1880-1933, Austrian), happened with the Rayleigh-Jeans law also proposed in 1900, which was trying to predict the shape of the light spectrum with the black body temperature. It was diverging to infinite values as the temperature was growing. Planck's law solved the problem and avoided the ultraviolet catastrophe. He found his spectrum equation empirically and only then, a related explanation based on harmonic oscillators and energy quanta exchanged between the radiation and the black body "wall". For this work, Max Planck was awarded the Nobel prize in Physics in 1918.

We also owe him the constant which bears his name (h) and which is used in his blackbody radiation equation. The Planck constant (6.626×10^{-34} Js) was then used in the equation according to which atomic state energy shift equals the radiation frequency multiplied by Planck's constant. The Planck constant is used in most




Figure 35: Quantum physics foundational years timeline. In green, experimentalist works, in black, theoretician works. The gold coins represent a Nobel prize. © Olivier Ezratty, 2021-2026.




Figure 36: Black-body spectrum and the ultra-violet catastrophe. The chart presents the Rayleigh-Jeans law that was created just before Planck’s law. It led to the ultraviolet catastrophe, with the law predicting an infinite energy level at small wavelengths. Planck’s law removed this “catastrophe” with his law which has zero spectral radiance flux at limit conditions. The peak temperature with the Planck law is independent of the composition of the black body. © Olivier Ezratty, 2022-2025.

quantum physics equations (De Broglie, Schrödinger, Dirac, etc.), often as the $\hbar$ (“hbar”) reduced version introduced by Paul Dirac, equal to $h/2\pi$.

When an electron changes its orbit in a hydrogen atom, it emits or absorbs an electromagnetic wave whose energy is equal to Planck’s constant multiplied by the emitted light frequency. More generally, a system can evolve only with multiples of Planck’s constant. Despite the numerous experimental validations carried out a few years later, Max Planck expressed until his death a lot of doubts about the very principles of quantum physics!

Planck is also at the origin of several infinitesimal or very large constants shown in Figure 37: Planck time, which is $t_P=10^{-44}$ s and Planck distance which is $l_P=1.616255 \times 10^{-35}$ m. Planck’s time is the time it would take for a photon to travel the Planck distance. These are the dimensions of the infinitely small below which any observation would be impossible. The length of Planck l_P is so small that a photon used to observe it would have such high frequency and energy that it would generate a black hole around it and would therefore become unobservable!

The Planck mass, about $22 \mu\text{g}$, is the mass scale at which the gravitational interaction between two particles becomes comparable to the other fundamental forces. It is loosely described as an upper bound for the mass of an elementary particle, since a heavier point-like object would collapse into a black hole. A particle with this mass and the size of Planck's distance would be a black hole. These are quite extreme physics. In today's classical cosmology, Planck's wall corresponds in the history of the expansion of the Universe to the moment when 10^{-43} s after the Big Bang, its size would have been 10^{-35} m, which is respectively the Planck time and Planck distance. Big Bang experimental conditions are difficult to reproduce. It doesn't prevent some physicists from trying to simulate it digitally [77].

At the extreme Planck temperature, the energy of particles is so high that gravitational forces between them become comparable to the other fundamental forces (electromagnetic, strong nuclear, and weak nuclear). It may have been the temperature at play at the beginning of the Big Bang. This is the absolute hot reference, in relationship with 0K, which is the absolute cold reference. Planck proposed these constants in 1899 just before he created his law explaining the black body radiation spectrum.

When participating in the second DPG Fall Meeting in Gottingen in September 2025, I discovered that Max Planck is buried in the city where Werner Heisenberg, Pascual Jordan and Max Born laid the foundations of modern quantum physics in 1925. But this is unrelated. He lived most of the war in Berlin. His home was bombed in 1944 and he ended his life in this city after world war II since he had relatives there.



Albert Einstein (1879-1955, German then American) got his Nobel prize in physics in 1921 for his interpretation of the photoelectric effect in 1905, which became one of the foundations of quantum mechanics after Planck and before De Broglie, Heisenberg and Schrödinger.

Einstein determined that Planck's quanta were elementary grains of energy $E = h\nu$ (Planck constant $\times$ frequency) with a momentum of $p = h\nu/c$ [78]. These were named "photons" in 1926 by **Gilbert Lewis** (1875-1946, American). The term was contested by some physicists like W.E. Lamb in his 1995 "Anti-photon" paper, because it emphasized the notion of particle at the expense of the notion of quantum fields and modes [79]. Symbolically, 1905 is also the year of Jules Verne's death.

Symmetrically to what Louis de Broglie would later do with electrons, he hypothesized that a photon behaves both as a wave and as a particle.

This was coming out of just one out of his four 1905 "annus mirabilis" papers sent to Annalen der Physik, the others being on Brownian motion in May, special relativity in June and mass-energy equivalence in September,

published when he was just 26. This was on top of his own 24 pages PhD thesis on a theoretical method to calculate molecular sizes using fluid mechanics and hydrodynamics.

With the photoelectric paper, he reconciled the corpuscular theories of **René Descartes** (1596-1650, French, in 1633) and **Isaac Newton** (1642-1726, English, in 1704) with the wave-based theories of **Christiaan Huygens** (1629-1695, Dutch, in 1678) to describe light. This was followed by the works from **Augustin-Jean Fresnel** (1788-1827, French), **Hippolyte Fizeau** (1819-1896, French, who made the first terrestrial measurement of the speed of light in 1849 and who gave his name to the Doppler-Fizeau effect), **Léon Foucault** (1819-1868, French, who showed in 1850 that light travels slower in water than in air, an argument in favor of the wave theory) and of course **James Clerk Maxwell** and this formalism describing electromagnetic waves.

The photoelectric effect corresponds to the capacity of a photon to dislodge an electron from an atom and to create some electric current (Figure 38), from a weakly bound valence electron at visible or ultraviolet frequencies up to an inner shell electron with X-rays. The electron layers of the atoms are numbered from 1 to N , their quantum number (or K, L, M, ...). One convention starts this labeling by K (first layer close to the nucleus with a maximum of 2 electrons) then L (8 electrons maximum), M (with a maximum of 18 electrons but in practice 8 for atoms other than transition metals due to different energy levels between layers 3/M and 4/N), etc. The photoelectric effect mainly concerns layers K and L.




Figure 38: The photoelectric effect, where a sufficiently energetic photon in the green or UV wavelengths hits the electron in a potassium structure, dislodging an electron.

The ejected electron is then replaced by an electron of external orbit, which generates a new photon, in X-rays or in fluorescence, according to the energy of the incident photon. This then emits an X-ray photon due to the energy differential between electronic layers or an electron called "Auger" from the name of Pierre Auger. This phenomenon was discovered around 1923 by the latter and by Lise Meitner. Another variant of the photoelectric effect is the Compton effect, when the high energy of an incident photon in gamma rays will release an electron from the valence layer and generate another photon. Finally, when the energy of the incident photon is even higher, the interaction takes place at the




Figure 37: Planck time, distance, mass and temperature constants. $\hbar$ is the reduced Planck constant introduced later by Paul Dirac, G is the gravitational constant, c is the speed of light and k_B is the Boltzmann constant. $\hbar = h/(2\pi)$ is a constant that was introduced after Planck created his constants. It is used in equations involving angular momentum and wave phenomena because of the natural appearance of 2π in angular frequency and wavenumbers. Planck time and mass has some consequences on the ultimate processing speed limit. Considering the capability to operate one bare physical operation per Planck time per Planck mass, we'd have the capability to support 10^{48} bare physical operations per second with one gram, corresponding to the processing speed of a black hole, as highlighted by Seth Lloyd in [24]. © Olivier Ezratty, 2021-2025.

nucleus of the target atom and generates an electron and a positron.

Silicon photovoltaic solar panel cells exploit a close cousin, the internal photoelectric effect, where the absorbed photon promotes an electron across the band gap without ejecting it out of the material. Photosynthesis rests on the same photon absorption and charge separation principle, which is the metabolic starting point of glucose production, although its detailed mechanism is a chain of molecular excitations and electron transfers rather than a photoelectric emission.

In addition to Max Planck's work on black body radiation, Einstein's interpretation was based on the earlier work of **Heinrich Hertz** (1857-1894, German) who discovered in 1887 that light can extract an electron out of metal, and **Philipp Lenard** (1862-1947, German) who, in 1902, studied the photoelectric effect and determined that it is only triggered at a certain frequency for the projected light. The latter was awarded the Nobel prize in Physics in 1905. Becoming a fervent Nazi and opposed to Einstein by scientific rivalry and then by explicit antisemitism, he mostly disappeared from the quantum physics hall of fame.

Einstein's photoelectric effect equations were then verified by the experiments of **Robert Andrews Millikan** (1868-1953, American) between 1914 and 1916, which provided one of the first direct measurements of Planck's constant. His separate oil-drop experiments, run between 1909 and 1913, measured the electric charge of a single electron. Both works earned him the Nobel prize in Physics in 1923.

Of course, Einstein is also at the origin of the special and general theories of relativity. He didn't obtain a Nobel Prize for his work on relativity despite its considerable impact on science. This is due, among other things, to his theories being based on earlier work from **Hendrik Antoon Lorentz** and **Henri Poincaré** as well as the contribution of his former professor **Hermann Minkowski** (1864-1909, German) who created the four-dimensional space-time notion in 1908 [80].

On top of many other contributions in quantum physics, Einstein predicted the photons stimulated emission effect in 1917, that would later lead to the creation of lasers. He also predicted in 1925 a particular behavior of matter, the Bose-Einstein condensate, which occurs when gases are cooled to very low temperatures. Atoms are then in a minimum energy quantum state showing particular physical properties. Superfluid helium-4, discovered in 1938 by Pyotr Kapitsa on one side and John Allen and Donald Misener on the other, is a related but not identical phenomenon since the strong interactions in the liquid mean that only about 10% of the atoms sit in the condensate. The first genuine dilute gas Bose-Einstein condensate came only in 1995, with rubidium at JILA by Eric Cornell and Carl Wieman and with sodium at MIT by Wolfgang Ketterle, which earned them the 2001 Nobel prize in physics [81, 82]. Bose is the name of **Satyendra Nath Bose** (1894-1974, India) with whom Einstein had worked during the 1920s and to whom we owe the "bosons", which verify the characteristics of Bose-Einstein's condensate statistics [83].

Bosons include elementary particles without mass such as photons and gluons but also certain composite objects such as the deuteron, which is the deuterium nucleus, and the helium-4 atom, as well as certain quasiparticles such as the superconducting electron pairs that are Cooper's pairs. We will see a little later that the spin sum of these particles determines whether they are bosons as opposed to fermions.

Albert Einstein also contributed to the philosophical-scientific debates on quantum physics realism, confronting Niels Bohr. He focused on the fact that quantum physics did not seem to completely describe the physical world with its probabilistic bias. Einstein wanted to find a realistic interpretation of quantum physics. He could not be satisfied with a probabilistic description of the state of electrons and other quantum objects. He could not find sufficient the interpretation of quantum physics according to which the observer and the measurement "make" the real world. He thought that the real world exists independently of measurements and observers.

The debate between Albert Einstein and Niels Bohr revolved around various thought experiments on determinism discussed during the 1927 Solvay Congress. It culminated later, in 1935, with the famous **EPR paradox** paper (Figure 40), named after its authors Albert Einstein, Boris Podolsky and Nathan Rosen. The paper raised the question of the incompleteness of quantum mechanics at the time [84]. The real author of the paper was Boris Podolsky and Albert Einstein thought it didn't really express his own views. While the thought experiment in EPR's paper didn't explicitly mention entanglement, a term that was created later the same year by Erwin Schrodinger, it related to two particles A and B that could interact for some time and not after it, that can now be considered as being entangled particles having some common past.

It sought to explain the nonlocality of the correlated quantum state measurement results of entangled particles which was a consequence of Schrödinger's wave function. It was not yet physically observed as of 1935. Einstein's view was that classical and relativistic physics act locally.

Gravity is local and is transmitted at the speed of light. All physical theories before quantum physics were local or EPR-local. Remote actions all involve a delay, usually coupled with attenuation with distance as it is the case for gravity.

For the EPR paradox paper authors, the quantum theory based on Schrödinger's wave function was either incomplete or two quanta could not be instantaneously synchronized at a distance at measurement time. Their measurement outcome being random and correlated, entangled quantum objects had to convey with them a sort of "information switch" indicating where the random measurement should land.

A physical theory is complete if each component of reality has a counterpart in the theory that makes it possible to predict its behavior, such as some tuning happening at the source when entangled quanta are created, and transmitted to each one with some hidden variables that would determine the outcome of their measurement. This underlies the notion of determinism, a principle that is absent in Schrödinger's wave function which is entirely probabilistic in nature.



Figure 39: The New York Times coverage of the EPR paper on May 4th, 1935, which infuriated Albert Einstein! The 1935 New York Times article was published thanks to a "leak" provoked by Boris Podolsky, the youngest of the EPR 3 gang. Source: [85].

The explanation of entanglement by "hidden variables" comes rather from Louis de Broglie with his pilot wave hypothesis elaborated in 1927, an idea later pursued by David Bohm in the 1950s [86]. With his "inequalities", John Stewart Bell demonstrated in 1964 that the existence of such hidden local variables was incompatible with the principles of quantum mechanics. Alain Aspect et al. 1982 experiments on photon entanglement did violate Bell's inequalities, which rules out local hidden variable theories, not quantum mechanics itself. In the end, Einstein could not finish his work on his theory of general relativity which was, for him, as incomplete as quantum mechanics. In particular, he wanted to reconcile quantum mechanics and gravity.

Einstein thought that quantum physics was an incomplete theory that didn't describe reality precisely enough. Einstein was then often credited with the idea that there were hidden variables. It seems, however, that he never mentioned them in his writings despite what John Stewart Bell later said. The EPR paper ends with indicating that it should be possible to build a complete theory of quantum mechanics (Figure 39). Hidden variables are a consequence rather than a hypothesis in the EPR paradox paper.

Be careful with the simplistic views that Einstein was "against" quantum mechanics, had it all wrong or did not believe in it [87]. He first questioned the principle of indeterminacy in 1927 and 1930, then estimated that the theory was incomplete to explain entanglement, with the EPR paradox paper, and finally, he opposed the lack of realism of quantum theory. This incompleteness is still being discussed more than 80 years later.

Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?

A. EINSTEIN, B. PODOLSKY AND N. ROSEN, *Institute for Advanced Study, Princeton, New Jersey*

(Received March 25, 1935)

In a complete theory there is an element corresponding to each element of reality. A sufficient condition for the reality of a physical quantity is the possibility of predicting it with certainty, without disturbing the system. In quantum mechanics in the case of two physical quantities described by non-commuting operators, the knowledge of one precludes the knowledge of the other. Then either (1) the description of reality given by the wave function in

quantum mechanics is not complete or (2) these two quantities cannot have simultaneous reality. Consideration of the problem of making predictions concerning a system on the basis of measurements made on another system that had previously interacted with it leads to the result that if (1) is false then (2) is also false. One is thus led to conclude that the description of reality as given by a wave function is not complete.

Figure 40: The famous EPR paper from Albert Einstein, Boris Podolsky and Nathan Rosen published in 1935 [84].

The origins of entanglement and nonlocality are still not physically explained. It is observed physically and described mathematically. This remains an open debate as scientists continue to ponder the different possible interpretations of quantum physics.

This is part of the field of quantum foundations that we cover later in this book.



Niels Bohr (1885-1962, Danish) was a physicist, Nobel prize in Physics in 1922, who created in 1913 a descriptive model of the hydrogen atom with its nucleus made of a proton and an electron rotating around the nucleus on precise orbits.

On these orbits, it is the orbital angular momentum of the electron that is quantized in multiples of $\hbar = h/2\pi$, h being Planck's constant and $n = 1, 2, 3$ and so on. The associated energy levels then scale as $-13.6/n^2$ eV for hydrogen. This model explained hydrogen spectral lines observed in the experiments of **Johann Balmer** (1825-1898) in 1885, **Theodore Lyman** (1874-1954) in 1906 and **Friedrich Paschen** (1865-1947) in 1908 (Figure 41, right). It also explained why electrons didn't crash on atom nucleus! Niels Bohr followed the work of **Ernest Rutherford** (1871-1937) who discovered in 1911 the structure of atoms with their positively charged nucleus, thanks to its protons, and their electrons revolving around the nucleus. The latter, with whom Niels Bohr worked in Manchester in 1912 after a first post-doc stay with J.J. Thomson in Cambridge in 1911, relied himself on **Hantaro Nagaoka** (1865-1950, Japan) who predicted in 1903 the structure of atoms with a positively charged nucleus and negatively charged electrons revolving around it, called the "Saturnian model".

Electrons had been discovered by **Joseph John Thomson** (1856-1940, English) in 1897 by analyzing the rays emitted by a cathode in a cathode ray tube (CRT), deflected by an electric field as well as by a magnetic field, and detected by a layer of phosphorus. He was awarded the Nobel prize in Physics in 1906.

Ernest Rutherford had also imagined the existence of neutrons, which was not verified experimentally until 1932 by **James Chadwick** (1891-1974, English). **Marie Curie** (1867-1934, Polish and French) had discovered polonium and radium in 1898 and some effects of radioactivity but not the existence of neutrons. According to Niels Bohr, electrons emit or absorb a photon when they change orbit. Subsequently, Louis de Broglie's work on wave-particle duality interpreted that the orbits of the electrons were an integer multiple of their associated wavelength.

Together with Werner Heisenberg, Pascual Jordan and Max Born, Niels Bohr is at the origin of the so-called **Copenhagen** interpretation of quantum physics which is based on three key principles:

- The description of a wave-particle is realized by its wave function, and no other "hidden" local information or variable can be used to describe its state. We must accept the wave function probabilistic nature used to describe a quantum state.
- When a quantum state measurement is performed, its composite wave function of several states is reduced to the wave function of one of the possible states of the quantum with a probability defined by Born's rule. This is the collapse of the wave function.
- When two properties are linked by an uncertainty relationship, the two properties cannot be measured with a greater precision than that allowed by the uncertainty relationship (Heisenberg principle of indeterminacy). Moreover, when we measure the position of a particle, we affect its motion, and vice versa. It comes from the bare fact that speed and position do not have any meaning before measurement in quantum physics. Variables linked through an indetermination link are conjugated with regards to actions which can change only by quantum leaps.

This is the main interpretation of quantum mechanics. There are many other interpretations. We will have the




Figure 41: The Bohr atomic model. Source: Wikipedia and other open-sources. 2023.

opportunity to detail the Copenhagen interpretation in the part on quantum foundations.

Note that Niels Bohr's son, **Aage Niels Bohr** (1922-2009, Danish), was awarded the Nobel prize in Physics in 1975 for his work on the structure of atom nucleus [88]!



Emmy Noether (1882-1935, German) created a theorem in 1915 at the University of Göttingen in Germany which says that if a system has a continuous symmetry property, then there are corresponding quantities whose values are conserved in time [89].

She created a second important and more general theorem that is the basis of gauge fields theories in quantum fields theory. At the origin of the field of abstract algebra, it is a foundation to Lagrangian mechanics, precursor of Hamilton's formalism. The equation below shows the conserved Noether quantity associated with a continuous symmetry of the action.

$$\frac{d}{dt} \left(\sum_a \frac{\partial L}{\partial \dot{q}_a} \delta q_a \right) = 0$$

At that time, she could not teach at the University because this role was forbidden to women. Her theorem was only published in 1918, and she could not officially teach until 1919. She did not receive a salary from the University until 1923. Her theorem links conservation principles and symmetries. It is one of the foundations of particle physics. Her work helped Albert Einstein to refine the foundations of the theory of general relativity he developed in 1915 [90]. She died relatively young, at 53.



Arthur Holly Compton (1892-1962, American) was a physicist who got the 1927 Nobel prize in Physics for the discovery in 1922/1923 of the effect which demonstrates that photons can have momentum and behave as particles.

His experiment makes a photon interact with a free electron around an atom, validating the photoelectric effect theories of Planck and Einstein (Figure 42). The Compton effect is a variant of this effect, applied to X and gamma rays which are high energy photons.




Figure 42: Compton scattering phenomenon. m_0 is the rest mass of the electron and c is the speed of light.

Compton scattering deals with the reception of an X or gamma photon whose energy is much higher than the binding energy of the target electron. The photon is deflected and comes out with a lower energy and a longer wavelength, a photon always travelling at c and never being slowed down. The photon-electron collision conserves energy and momentum like an elastic two-body shock, but since the photon wavelength changes, the process is classified as inelastic or incoherent scattering.

The Compton effect shows up in X-ray radiography where, unlike the photoelectric effect which creates the image contrast, it mostly generates scattered radiation that degrades it. X-rays are emitted during electronic transitions between the atomic layers K, L and M (the first around the nucleus of the atom). The emission angles of the ejected electron and the re-emitted photon depend on the incident photon energy level.



Jacques Salomon Hadamard (1865-1963, French) was a mathematician who worked on complex numbers, differential geometry and partial differential equations (PDEs), particularly during the 1920s.

He also became interested in the creative process of mathematicians by studying the creative process of hundreds of colleagues. His name was given to the Hadamard single qubit gate used in quantum computers and quantum algorithms which creates a superposed state between $|0\rangle$ and $|1\rangle$. It is a special case of the Walsh-Hadamard transforms, square matrix operations with 2^n entries on each side, whose recursive construction actually goes back to James Joseph Sylvester in 1867, Hadamard having later characterized in 1893 the maximal determinant matrices that bear his name, as described in Figure 43. The single qubit quantum gate named after Hadamard is a transform of Hadamard of type H_1 . It switches the Z basis for the X basis around the X+Z axis of the Bloch sphere (Figure 201). This superposition is one of the enablers of computing parallelism in quantum computing, in addition to the principle of entanglement which links the qubits together and is one of the sources of quantum exponential acceleration. Superposition is only responsible for a potential polynomial acceleration.

$$H_0 = +(1)$$

$$H_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

$$H_2 = \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}$$

$$H_3 = \frac{1}{2^{3/2}} \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & -1 & 1 & -1 & -1 & 1 & -1 & 1 \\ 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \\ 1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 \end{pmatrix}$$

Figure 43: Walsh-Hadamard transforms for 0, 1, 2, and 3 qubits.



Otto Stern (1888-1969, German-American) and **Walther Gerlach** (1889-1979, German) respectively conceived in 1921 and together realized in 1922 in Frankfurt the famous Stern-Gerlach experiment.



The experiment helped discovering the intrinsic angular momentum (or spin) quantization in a magnetic field using a beam of electrically neutral silver atoms as shown in Figure 44.

The X, Y and Z components of the electron spin measured in the Stern-Gerlach experiment are complementary variables. Measuring one of the three variables prevents from doing so with the two others. In the experiment, this momentum came from the 47th electron' spin from heated silver atoms [91].

It showed that these atoms have a quantized angular dipole that deflects the beam in a given direction upward or downward. It later became known as particle spins a bit later, in 1925, according to the work of George Uhlenbeck and Samuel Goudsmit. The experiment also showed that spin measurement along a given direction was incompatible with being done in another direction, corresponding to the notion of observables complementarity.




Figure 44: The Stern-Gerlach experiment where an atomic stream of silver is deviated in two discrete directions by a magnetic field. Image source: Wikipedia. 2023.



Louis de Broglie (1892-1987, French) was a mathematician and physicist who, in 1923 and 1924, extended the particle-waves duality, then only applied to photons, to massive particles, mainly electrons, and atoms, protons and neutrons.

According to this principle, elementary particles behave like particles (with a position, a trajectory and possibly a mass) and like waves (potentially delocalized and scattering in all directions and generating interference) depending on the circumstances.

$$\begin{array}{ll} \text{particle energy} & \text{---} \bullet \text{---} E = h\nu \quad \bullet \text{---} \text{ wave frequency} \\ \text{particle momentum} & \text{---} \bullet \text{---} p = \frac{h}{\lambda} \quad \bullet \text{---} \text{ Planck constant} \\ & \quad \quad \quad \bullet \text{---} \text{ wave length} \end{array}$$

Figure 45: De Broglie wave-particle equation with electrons.

This is the case of electrons which have a mass and can interfere with each other. Louis de Broglie turned this duality into an equation: $\lambda p = h$, where λ is a wavelength, p is a quantity of motion and h is Planck's constant (Figure 45).

This earned him the Nobel prize in Physics in 1929. He is the main French contributor to quantum physics during the inter-war period. Louis de Broglie's brother,

Maurice de Broglie (1875-1960), was also a physicist. He had studied X-rays and spectrography. Both brothers were members of the Academy of Sciences in France.




Figure 46: Electron wave-particle diffraction experiment. Source: [92].




Figure 47: Electron wave-particle diffraction experiment details. Source: [92].

The wave-particle duality of electrons was confirmed in 1927 as shown in Figure 46 and Figure 47 with a nickel crystal based diffraction experiment by **Clinton Davisson** (1881-1958) and **Lester Germer** (1896-1971) from the Bell Labs in the USA, who shared a Nobel prize in physics in 1937.

The electron wavelength is in the picometer range and is much smaller than visible photon wavelengths. It explains why electron microscopes have a better resolution than classical optical microscopes.

George Paget Thomson (1892-1975) from the University of Aberdeen in Scotland did a similar experiment also in 1927. However, the Young double-slit experiment done with electrons was realized much later, in 1961, by **Claus Jönsson** (1939, German).

The wave nature of neutrons was already established in 1936 with crystal diffraction, by Hans von Halban and Peter Preiswerk in Paris and by Dana Mitchell and Philip Powers in the USA, and neutron interferometry followed in 1974 with Helmut Rauch in Vienna. The

single and double slit version came only in 1988 with **Anton Zeilinger**, **Roland Gähler** et al. [93]. Atom diffraction was demonstrated in 1991 by **Olivier Carnal** and **Jürgen Mlynek**, using double-slit diffraction and by **Mark Kasevich** and **Steven Chu**, who created the first cold atom interferometer using a light-beam splitter based on Raman transitions. It became the basis of atom interferometry used in quantum absolute gravimeters, using an equivalent of a Mach-Zehnder interferometer replacing light with matter-wave made of atoms. In this setup, the Mach-Zehnder beamsplitter is replaced by a series of three lasers pulses creating a superposition of two atomic energy states driving a diffraction effect, then a mirror effect and at last for a recombination of split wave packets. It is even verifiable with molecules of several atoms.



Wolfgang Pauli (1900-1958, Austrian, then American and Swiss) is at the origin of the principle of exclusion which bears his name elaborated in 1925 and according to which two electrons cannot have the same quantum state in an atom.

He first discovered in 1924 the atom nucleus spin, used to explain the hyperfine structure of atomic spectra, i.e., the existence of very close spectral lines observed during their excitation. It cannot be explained by quanta and energy levels of the electron layers in the atoms.

In 1925, he formulated the exclusion principle according to which electrons in the same system (an atom) cannot be simultaneously in the same quantum state, a principle that was later extended to all fermions, i.e., half-integer spin particles (electrons have a spin $\frac{1}{2}$ but fermion atoms can have $\frac{3}{2}$, $\frac{5}{2}$, $\frac{7}{2}$ and even $\frac{9}{2}$ spins, like ^{40}K). For example, the two electrons in a helium atom are in the same shell and must have an opposite spin. On top of that they are entangled since they are indissociable.

He then proposed between 1925 and 1927 that the electron has an additional theoretical degree of liberty, on top of the first three quantum numbers describing the state of an electron in an atom, with the mathematical formalism to describe it using the famous Pauli matrices and operators. The first electron quantum number is the energy level of the electron in the atom (the layer where it is located), the second is the azimuthal quantum number or orbital quantum number (which defines the electron sub-shell) and the third is the magnetic quantum number (which describes a discretized orientation of the subshell and makes it possible to distinguish the orbitals of the electron in the atom).

The second and third electron quantum numbers were introduced by Arnold Sommerfeld (1868-1951, German). Among others, Wolfgang Pauli and Werner Heisenberg were his PhD students.

The alpha constant or fine structure constant is also called the Sommerfeld constant per his work from 1916 [94].

The fourth degree of freedom of the electron was identified in 1925 by **George Uhlenbeck** (1900-1988, The Netherlands and USA) and **Samuel Goudsmit** (1902-1978, Dutch/American) as an intrinsic angular momentum or electron spin. George Uhlenbeck and Samuel Goudsmit were students of Paul Ehrenfest (1880-1933, Austria/the Netherlands). His laboratory had welcomed some illustrious future physicists such as Enrico Fermi, Robert Oppenheimer, Werner Heisenberg and Paul Dirac. Ehrenfest was a specialist in statistical physics. In particular, he contributed to the understanding of phase changes in matter. His theorem formulated in 1927, creates a direct mathematical bridge between quantum dynamics and classical Newtonian equations of motion. It states that the quantum expectation values of position and momentum obey equations that formally mirror classical mechanics.

We currently do not have an image or physical representation of what the spin is [95]. The argument that a spinning electron with a classical radius would have a surface velocity exceeding the speed of light is usually credited to Hendrik Lorentz, who raised it in 1925 when Uhlenbeck and Goudsmit submitted their result, Pauli having separately objected to Ralph Kronig's earlier version of the spin idea. Electron spins are used in silicon qubits that we cover later. He also conjectured the existence of neutrinos in 1930, which was experimentally proven in 1956, and worked on quantum electrodynamics. He was awarded the Nobel prize in Physics in 1945.

Otherwise, 137 is a number that played a weird role in Pauli's life. It turns out that $1/137$ is a value that roughly corresponds to the fine-structure constant, a dimensionless ratio that is found in many places in quantum physics that compares data with the same dimension. The fine-structure constant was measured with a relative accuracy of 2.0×10^{-10} in 2018 by Richard Parker, Holger Müller et al. at Berkeley using cesium atom interferometry [96], and then to 8.1×10^{-11} in 2020 by Léo Morel, Saïda Guellati-Khélifa et al. at LKB in Paris using rubidium [97]. It is for example the ratio between the velocity of an electron in the lower layer of a hydrogen atom and the speed of light or the probability of emission of the absorption of a photon for an electron [98]. "137" is a sort of "42" of quantum physics. Wolfgang Pauli died after some pancreatic cancer surgery, while his hospital room number was 137!



Werner Heisenberg (1901-1976, German) was a physicist who created with Max Born and Pascual Jordan in 1925 the quantum matrix formalism describing physical quantities. He was awarded the Nobel prize in Physics in 1932.

In 1927, he also created the famous principle of uncertainty, or rather indeterminacy, according to which one cannot accurately measure both the position and the velocity of an elementary particle, or, more generally, two arbitrary unrelated quantities.

The indeterminacy principle is a consequence of this formalism. It was described mathematically in a simplified manner in 1927 by **Earle Hesse Kennard** (1885-1968, American) in the famous equation below, where the product of the standard deviations of position and momentum is greater than half the Dirac (or reduced Planck) constant.

$$\Delta x \Delta p \geq \frac{\hbar}{2}$$

This principle is used to improve the accuracy of a measurement of any quantity by lowering the accuracy of another quantity characterizing a quantum state. This technique is used in "quantum squeezing" which is integrated into the latest version of LIGO for the measurement of gravitational waves [99]. These quantities can be for example an energy level, a position, a wavelength, or a speed. One consequence of Heisenberg's indeterminacy principle is that all particles in the Universe are in permanent motion. If they were stable, we would know their position (fixed) and their velocity (zero), violating the indeterminacy principle.

Another consequence is that a perfect vacuum could not exist because the value and evolution of the magnetic and gravitational fields that pass through it would be stable, violating once again Heisenberg's indeterminacy. This explains the astonishing vacuum quantum fluctuations that we will discover a little further down this book. The no-cloning theorem of a qubit state also derives from the principle of indeterminacy.

For some, this indeterminacy principle is a simplified interpretation of the corpuscular nature of matter. It leads to the question of the position and velocity of an electron when it has no precise position. According to the Copenhagen interpretation of quantum mechanics, we shouldn't try to determine where the electron is located.




Figure 48: Heisenberg microscope thought experiment. [Source](#).

In practice, quantum particles are not classical physical particles and therefore their velocity and position cannot be measured. They can only be described by their

(Schrödinger) wave function and position probabilities. More generally, in the infinitely small, the measurement device influences the measured quantity. One example illustrates this phenomenon at the macroscopic level: if you illuminate an insect with sunlight and a magnifying glass to better observe it, you may burn it! The same happens with a photon that is used to detect an electron, in the Heisenberg microscope thought experiment, as shown in Figure 48. It will change the speed and position of the electron.

Like many of the colleagues of his time, Werner Heisenberg was interested in the links between science, quantum mechanics and philosophy, and as early as 1919. He was an assistant to Niels Bohr between 1924 and 1927, before leaving for the University of Leipzig. Max Born was also one of his professors.

During World War II, he was asked with other German scientists to work on the Reich's atomic bomb project. Later revelations did show that he was not very active on this project and did not believe it was an achievable goal. He even discussed it with Niels Bohr during the war, in September 1941 in Copenhagen, which was then occupied by the Germans [100]. It even led to the creation of the play "Copenhagen" by Michael Frayn which was published in the UK in 1998.



Erwin Schrödinger (1887-1961, Austrian) is a physicist who was awarded the Nobel Prize in 1933 for the creation of his famous wave function in 1926, *aka* Schrödinger equation.

It describes the evolution in time and space of the quantum state of a massive quantum particle and the probabilities of finding the quantum at a given place and time. Schrödinger's equation is a variant of the Newtonian mechanics equations that define the total energy of an object as the sum of its kinetic energy and its potential energy. We describe this equation in detail in the Quantum 101 part of this book (3.3.1).

Erwin Schrödinger also created his famous alive and dead cat in a box thought experiment. The cat thought experiment was published in a series of three papers in 1935, shortly after the publication of the EPR paradox paper by Einstein, Podolsky and Rosen [101]. The cat story fits in only about ten lines in this long document which deals with superposition, measurement, and entanglement. In the paper, Schrödinger also coined the term entanglement in the first chapter "*The Lifting of Entanglement. The Result Depends on the Will of the Experimenter*". Schrödinger translated himself the German word "*Verschränkung*" into "*entanglement*", which he introduced in English in his 1935 Cambridge Philosophical Society paper. The cat that appears only three times in all and for all is therefore anecdotal but that is what everyone has remembered. The story was hidden for a while and revived in the early 1980s, particularly after Alain Aspect's experiment [102].

In Schrödinger's original scenario, a steel chamber contains a flask of prussic acid, the breaking of which is triggered by the decay of one atom from a tiny amount of radioactive substance, detected by a Geiger counter. Schrödinger names no particular isotope, and the quantity is chosen so that a single atom has about one chance in two of decaying within the hour, which gives the cat a 50/50 chance of being alive or dead when the hour is up. When opened, it is either alive or dead.

As long as the door is not opened, the cat is said to be superposed in the alive and dead states and entangled with the radium atom state. This story has been repeated ad-nauseam since 1935. But his thought experiment was created to show the absurdity of the measurement postulate, the wave function collapse and Born's rule. Unfortunately, the contrary has been memorized (Figure 49).




Figure 49: The infamous Schrodinger's cat thought experiment raw description.

The caveat is that a cat can't be superposed in two states because it is a macroscopic object of a size well beyond the quantum/classical limit. It is either alive or dead, never both. These are exclusive states. They can't be implemented as linear superpositions of waves. When the cat is dead, he's not in a superposition. He's just plain dead [103]. On top of that, the radium atom disintegration as well as the cat's death are both irreversible processes.

But Ronald Gurney and Edward Condon demonstrated in 1928, before the Schrodinger's paper, that nuclear decay is governed by a wavefunction with amplitudes describing both "before decay" and "after decay" configurations [104]. As a result, the state of the nucleus can in principle be a coherent combination of "nucleus intact" and "nucleus decayed" components until an interaction with the environment like the emitted particle's detection makes one branch effectively irreversible. This is hard to visualize!

We can also consider that the cat's death is provoked by a not yet read measurement when the box is closed, corresponding to a non-selective measurement. The cat state uncertainty is a classical one, not a quantum one. The cat is in a maximally "mixed state" where the uncertainty of its death is classical, not in a "pure state" where it would be quantum. If you used a webcam

inside the box and made sure it didn't influence the radium half-life period, you could track the cat state all along, from alive to dead or alive to alive, which are the only two possible paths and observe the absence of superposition.

The Schrodinger text says exactly: *“It is also possible to construct very burlesque cases. Imagine a cat locked up in a room of steel together with the following hellish machine (which has to be secured from direct attack by the cat): A tiny amount of radioactive material is placed inside a Geiger counter, so tiny that during one hour perhaps one of its atoms decays, but equally likely none. If it does decay then the counter is triggered and activates, via a relais, a little hammer which breaks a container of prussic acid. After this system has been left alone for one hour, one can say that the cat is still alive provided no atom has decayed in the mean time. The first decay of an atom would have poisoned the cat. In terms of the ψ -function of the entire system this is expressed as a mixture of a living and a dead cat”.*

This thought experiment was intended to highlight two things. First, that superposition and entanglement only applied to the infinitely small and not to macroscopic objects. History retained the principle of superposition and not this difference between the microscopic and macroscopic worlds. Second, that there was and still is an uncertain limit between the quantum and classical worlds.

Schrodinger's thought experiment also dealt with the entanglement between the radium atom and the cat. Could this entanglement work with a macroscopic object? The paper containing this thought experiment was about entanglement and that was forgotten. Also, this paper's publication was the one generating the creation of the EPR paradox piece by Einstein et al. We should remember that Schrödinger's wave function and the notion of states superposition only make sense at a microscopic scale. Let's leave that poor cat alone in his dreams!



Max Born (1892-1970, German) is a physicist and mathematician who developed the mathematical representation of quantum in a matrix form along with Werner Heisenberg. He was awarded the Nobel prize in physics in 1954.

We owe him in 1926 the statistical explanation of the probability of finding an electron in a given energy state from its wave function, elaborated by Schrödinger the same year. This principle is applied to qubits, where the amplitudes attached to the two basis states are complex numbers and where the sum of their squared moduli is equal to 1. The probability of each outcome is the squared modulus of its amplitude, and not the amplitude itself.

In 1925, he created the non-commutativity relation of two conjugate quantities, one being the Fourier transform of the other (the commutator $[X, P] = XP - PX =$

$i\hbar I$, where X is a position and P a momentum and I , the identity). It led to the creation of the indeterminacy principle. Max Born also created the first version of the adiabatic theorem with Vladimir Fock in 1928.

Fun fact, the British singer Olivia Newton-John is his grand-daughter [105]. And he was also the PhD Director of Robert Oppenheimer, the future Director of the Manhattan project.



Paul Dirac (1902-1984, English) is a mathematician and physicist among the founders of 20th century quantum physics. He is credited with the 1928 electron spin equation, which is one of the foundations of relativistic quantum physics.

His equation is a kind of variant of Schrödinger's equation for free relativistic particles, fermions (electrons, protons, neutrons, quarks, neutrinos) which are half-integer spin particles. Relativistic particles are those moving at a speed close to the speed of light, like electrons in lower shells of heavy atoms.

In Dirac's equation, the wave function ψ of the electron includes four components of complex numbers that integrate time and space. Dirac's equation enabled him to predict the existence of a particle that was later called the positron, an opposite of the electron with a positive charge.

$$\left(\beta mc^2 + c \sum_{n=1}^3 \alpha_n p_n \right) \psi(x, t) = i\hbar \frac{\partial \psi(x, t)}{\partial t}$$

Positrons were discovered experimentally by Carl Anderson in 1932. He was awarded the Nobel prize in physics in 1936.

Dirac formalized the quantization of the free electromagnetic field in 1927. He also introduced in 1939 the bra-ket notation, known as Dirac's notation, which simplified the notation and manipulation of quantum states and operators in linear algebra (example: $\langle \phi | \psi \rangle$). The Dirac constant also named reduced Planck constant is the Planck constant h divided by 2π , also called “hbar” for its italicized strikethrough h symbol: $\hbar$. This Dirac constant is used in the Schrödinger wave function.

Paul Dirac was awarded the Nobel prize in Physics in 1933, at the age of 31. The Nobel Prizes of the early 20th century were frequently awarded to young scientists, which seems to be out of fashion since then! The youngest Nobel prize in physics was awarded to Lawrence Bragg, who won it at the age of 25 in 1915, shared with his father William Henry Bragg, for the analysis of crystal structures by X-ray diffraction, work he had started at the age of 22.

Paul Dirac was special for his shyness and parsimonious oral expression in meetings or during meals. So much so that his Cambridge colleagues had defined the “Dirac” unit as the most concise way to express himself in a

meeting, namely, at the rate of a single word per hour. His behavior was equivalent at the Solvay Congresses he attended, notably that of 1927. However, he broke a record in his Nobel Prize acceptance speech at the end of 1933 [106]. It is six pages long! Half, however, of the 12 pages of the speech of Erwin Schrödinger, also winner of the Nobel prize in physics that year. Another anecdote: Dirac was married to one of the sisters of Eugene Wigner, Nobel prize in physics in 1963 and famous for his function and also his “friend” paradox.

In which case do we have to deal with relativistic particles, in particular with electrons? It is generally considered that an electron becomes relativistic when the total of its mass and kinetic energy is at least twice the rest mass. This ratio corresponds to the **Lorentz factor**^w. It represents a speed of at least 86% of the speed of light (Figure 50). But relativistic phenomena may occur before that speed is reached. In Newtonian equivalent, the speed of an electron around the nucleus of a hydrogen atom is about $c/137$. With electrons from heavy atoms inner shells, this velocity can exceed $c/2$.

This affects the position of relativistic electrons in the low orbits of heavy atoms such as lanthanides, which belong to the rare earths. The Bohr radius that defines the average orbital of an electron decreases inversely proportional to the apparent mass of the electron. Because the electron’s apparent mass increases, this Bohr radius is smaller for relativistic electrons. This modifies the structure of the electron orbitals of heavy atoms and the transition energy levels between orbitals that absorb or emit photons.




Figure 50: Relativistic electrons and Lorentz factor. Relativistic atomic phenomena are said to start at 86% of the speed of light. They occur with inner layers electrons in relatively heavy elements.

This explains the color of gold and silver, due to relativistic modification of orbits of electron layers between which transitions occur due to the absorption of photons. Blue is absorbed in the case of gold, explaining its yellow color. Without the relativistic effect, gold would be white. This has a lot of implications in the chemistry of these materials and with their crystal organization [107]. This quantum relativistic effect also explains why

mercury is liquid at room temperature [108]. All this gives rise to a field of chemistry called relativistic quantum chemistry [109–111]. It also explains why the size of atoms is not proportional to their number of protons and electrons [112].

Particles also become relativistic in **particle accelerators** such as the CERN LHC near Geneva (the largest in the world), the ESRF in Grenoble (European Synchrotron Radiation Facility, specialized in the generation of “hard”, very high-frequency X-rays) or the SOLEIL light synchrotron located in Saint-Aubin near Saclay just next to the CEA, also in France, or its equivalent from PSI in Switzerland.

The SOLEIL synchrotron uses electrons accelerated to a relativistic speed and inverters that generate beams of light 10,000 times denser than sunlight [113]. Electrons circulate in the synchrotron at a speed close to the speed of light. SOLEIL powers more than 25 analytical instruments covering the spectrum from infrared to X-rays, with numerous applications in precision microscopy, including a microscopy using very well collimated and polarized white light. These instruments can be used to analyze the three-dimensional structure of organic molecules such as complex proteins, such as the glycoproteins that surround viruses. This even allows one to study how these proteins combine with those of the attacked cells, or ribosomes, which are used to produce proteins in cells. Equivalent instruments exist such as the Advanced Photon Source at the Argonne National Laboratory from the US Department of Energy near Chicago.

Free Electron Lasers (FEL) exploit relativistic electron sources. These are lasers generating coherent light (spatially and temporally, the emitted photons have the same frequency, phase and in that case, also polarization) and exploit relativistic electron sources from synchrotrons (Figure 51).

The interaction between these electrons and a strong alternating magnetic field makes it possible to generate coherent light in electromagnetic frequency ranges from infrared to X-rays, through visible light and ultraviolet. The FEL are used to explore all sorts of matter, particularly in biomedical research like with X-rays crystallography.

Finally, relativistic particles can be found in astrophysics and, for example, in cosmic ray sources as well as in relativistic plasma jets produced at the center of galaxies and quasars [115]. Dirac’s equation is linked to the **Klein-Gordon equation** (1926). The Klein-Gordon equation applies to spin-0 particles such as pions and the Higgs boson. Spin-1 gauge bosons like photons and gluons are described by the Proca and Yang-Mills equations instead. Relativistic quantum mechanics is a broad field of physics, used in elementary particles physics.




Figure 51: Free-electron laser. Source: [114].



Vladimir Fock (1898-1974, Russian) was a theoretician physicist who worked on quantum physics, the theory of gravitation and theoretical optics. We owe him the Fock space, representation and state.

It is used in quantum photonics to represent the state of bosons many-body systems having the same quantum state. He co-created the Klein-Gordon equation in 1926, the relativist version of Schrödinger's equation for zero spin massive particles, the adiabatic theorem with Max Born in 1928 and the Hartree-Fock quantum simulation method in 1930. He also worked on quantum electrodynamics and quantum foundations.



Pascual Jordan (1902-1980, German) was a physicist who collaborated with Max Born and Werner Heisenberg and contributed to laying the mathematical foundations of quantum mechanics, especially in matrix computation.

Like Philipp Lenard, he was somewhat forgotten because of his membership in the Nazi Party during the 1930s, although he was rehabilitated after the World War II thanks to the help of Wolfgang Pauli. He became interested in the philosophical notion of free will.



Linus Pauling (1901-1994, American) was a physical chemist known to have co-founded the scientific fields of quantum chemistry and molecular biology. He had the opportunity to meet in Europe the founders of quantum physics like Erwin Schrödinger and Niels Bohr in 1926-1927.

He described chemical bonds over a period between 1928 and 1932 and the hybridization of orbitals which explains the geometry of molecules. He published reference book “*The Nature of the Chemical Bond*” in 1939. He was awarded the Nobel Prize in Chemistry in 1954 and the Nobel Peace Prize in 1962 for his political activism in favor of nuclear disarmament. He is at the origin of computational chemistry, which makes it possible to numerically simulate the structure of molecules and which we discuss in the section on quantum applications in healthcare.



James Chadwick (1891-1974) is an English physicist who was responsible for the discovery of neutrons in 1932, which earned him the Nobel prize in Physics in 1935. This discovery was later compared to quantum physics and the discovery of electrons.

Nuclear physics has indeed progressed in parallel with quantum physics, which was mainly concerned with the interactions between electrons and photons. Before the discovery of neutrons, scientists thought that the nucleus of atoms contained protons and electrons.



John Von Neumann (1903-1957, Hungarian, then American) was a polymath and an extremely prolific mathematician. He contributed to the development of quantum mechanics, the atomic bomb, and classical computing.

He participated in the creation of the mathematical foundations of quantum mechanics, notably in the “*Mathematical Foundations of Quantum Mechanics*” published in 1932 [116]. He transposed the main principles of quantum mechanics into models and equations of linear

algebra. He devised the key mathematical principles behind quantum measurement models.

This deals, for example, with the representation of quantum states as a position in a Hilbert space, the observables which are projections into Hilbert spaces and the indeterminacy principle which can be explained by the non-commutativity of measurement operators. These principles are also named Birkhoff-von Neumann *quantum logic*, in connection with their seminal paper published in 1936 [117].

von Neumann also affirmed in 1932 that the introduction of hidden variables to incorporate determinism was a lost cause because it would contradict other (verified) predictions of quantum physics. Three years before Einstein/Podolsky/Rosen's EPR paper! His proof was however too restrictive, as Grete Hermann showed as early as 1935 and as John Bell demonstrated again in 1966: it silently assumed that expectation values add linearly for non-commuting observables, an assumption that no hidden variable theory has to satisfy [118].

We owe him the creation of the notion of entropy (by Von Neumann), in 1932, which is associated with the notions of operators and density matrices that he created in 1927 and which describe the state of a multi-partite quantum system. He participated in the Manhattan project in the USA where he modelled explosions and lenses for compressing plutonium in A-bombs.

He also created basic concepts in game theory and classical computers that are still in use. Almost all computers use a von Neumann architecture with memory, registers, control unit, computing unit, inputs and outputs as shown in Figure 52.

He designed it in June 1945 while contributing to the design of EDVAC, the first stored-program computer architecture. The first machine to actually run a stored program was the Manchester Baby in June 1948, EDVAC itself being delivered only in 1949. It became known as the Princeton architecture. When he wrote the First Draft of a Report on the EDVAC in 1945 [119], von Neumann was working with John Mauchly and J. Presper Eckert from the University of Pennsylvania.

However, he was formally affiliated with the Institute for Advanced Study in Princeton, which is why the architecture is sometimes referred to as the Princeton architecture. Its name was chosen in contrast with the then "Harvard architecture" which used separate memory and buses for instructions and data, while the Von Neumann-Princeton model used a single memory and bus for both instructions and data. Nowadays, generalist computers use primarily the Princeton model while embedded systems often use the Harvard model, but the architectures are becoming hybrid particularly for modern CPUs.

What a contribution!

PRINCETON (VON NEUMAN) ARCHITECTURE




Figure 52: The Von Neuman Princeton architecture which still defines classical computing. It was defined in 1945 in [119].



Boris Podolsky (1896-1966, Russian-American) wrote the EPR paradox paper with Albert Einstein and Nathan Rosen in 1935 on quantum entanglement and questions of nonlocality of the properties of entangled quanta.

He was a specialist in electrodynamics which deals with the analysis of electric and electromagnetic fields. He emigrated to the USA and, according to Soviet archives declassified in the 1990s, passed information on the American atomic program to Soviet intelligence, then the NKVD, around 1942-1943. His code name was... "Quantum".



Nathan Rosen (1909-1995, American then Israeli) is the third EPR paradox author when working as an assistant to Albert Einstein in Princeton. After moving to Israel in 1953, he created the Institute of Physics at Technion University in Haifa.

He was mainly working on astrophysics and relativity theory. He devised the concept of wormholes, a theoretical link between different points in space and time. He also thought neutrons were built out of a proton coupled to an electron.



Ettore Majorana (1906-~1938, Italian) imagined the existence of fermions in 1937 based on Dirac's equations, elementary particles that would be their own antiparticle.

His contributions also spanned atomic spectroscopy which led to the Thomas-Fermi model, atom nucleus models and neutrons, ferromagnetism, quasi-stationary states and quantum field theory. The Majorana fermion naming is also abusively applied in condensed matter physics to quasiparticles having similar properties.

Their existence was discovered in 2012 and verified in 2016 and then in 2018, even if it is still disputed by many physicists and two related 2018 papers had to

be retracted in 2021. These Majorana quasiparticles (or “Majorana Zero Modes”) could make it possible to design universal quantum computers called topological computers that can handle very efficient error correction codes requiring a small number of physical qubits. This is the exploration path chosen by Microsoft after the work of Michael Freedman and Alexei Kitaev in the late 1990s. Ettore Majorana is said to have committed suicide after a depression, because he could hardly stand the pressure of his genius! But his disappearance remains enigmatic because his body was never found!



Alonzo Church (1903-1995, American) was a mathematician who was a key contributor to the foundations of theoretical computer science and on the notion of computability. Among other things, he created the lambda calculus in 1936

It is a universal abstract programming language which inspired the creation of LISP. He also proposed the so-called Church-Turing thesis for which any automatic calculation can be carried out with a Turing machine. Church and Turing also proved an equivalence between being λ -computable and Turing computable.

Many variations of the Church-Turing thesis were elaborated after them to extend the broad field of complexity theories. For example, the extended Church-Turing thesis states that any physically realizable computing model can be simulated by a probabilistic Turing machine with at most a polynomial overhead. It is not demonstrable, and quantum computing is precisely the main candidate for falsifying it.

What about the others, known, unknown or less famous from the 1927 Solvay Congress? Two participants deserve to be mentioned who had some connections with quantum physics.



Léon Brillouin (1889-1969, Franco-American) who is less known in France because of his expatriation to the USA during World War II contributed to advances in quantum physics between the two World Wars.

In particular, he brought quantum mechanics closer to crystallography with the notion of Brillouin zones, which delimit the allowed electron energy bands in a periodic lattice. He separately predicted in 1922 Brillouin scattering, the inelastic scattering of light on the acoustic vibrations of a medium, which shifts the scattered photon frequency by a few GHz.

And then, finally, **Hendrik Anthony Kramers** (1894-1952, Dutch) who assisted Niels Bohr in the creation of quantum theory. Many of the participants were not quantum physics scientists. They were invited because the Belgium organizers tried to have a stable proportion of Belgians, French, Germans and English participants. Were there, for example, **Émile Henriot** and **Marie Curie** who were focused on radioactivity, way before it could be explained by the formalism of quantum physics,

Paul Langevin (with whom Marie Curie had had an affair in 1910, after the accidental death of her husband Pierre Curie in 1906), as well as a good number of chemists.

What was striking during this prolific period was the way the social network of physicists worked, without smartphones and the Internet. They had many encounters, cross-University tenures, meetings, letter exchanges and conferences. It was slow according to today’s references, but the results were still astounding.

To conclude this part, let’s remind us how young the founders of quantum physics were when they published their seminal work in the key years from 1900 to 1935 (Figure 53). Back then, scientific research didn’t work the same way. They also were frequently awarded Nobel prizes at less than 40! Nowadays, most of the times, you must wait until you are at least 50 if not 70.

2.3 Post-war

As mentioned before, quantum physics developments seemed to slow down between 1935 and 1960. Physicists were then busy with nuclear physics. The Manhattan project mobilized an amazingly large number of physicists like John von Neumann and **Enrico Fermi** (1901-1954, Italian American, Nobel prize in physics in 1938) whose contributions were centered in nuclear physics (first nuclear reactor in 1942). Pre-war, he was also a key contributor to quantum physics with his work on statistical physics, leading to the Fermi-Dirac ideal gas statistics, fermions and Fermi sea (1926), Fermi’s golden rule, actually derived by Paul Dirac in 1927 and named after Fermi who called it his “golden rule number 2” in his 1950 lecture notes and his work on neutrinos (1934).

This led to the creation of several research labs (Los Alamos, Oak Ridge, etc.) which are now DoE labs and host strong academic research organizations in quantum technologies.

Quantum physics still led, after World War II, to an incredible wealth of technologies that revolutionized the world (Figure 54). We can mention three important branches resulting from the applications of the first quantum revolution: **transistors**, with the point-contact transistor demonstrated in December 1947 by John Bardeen and Walter Brattain at the Bell Labs and the junction transistor invented by William Shockley in 1948 [120], **masers** and **lasers** invented between 1953 and 1960 by Gordon Gould, Theodore Maiman, Nikolay Basov, Alexander Prokhorov, Charles Hard Townes and Arthur Leonard Schawlow, only a few of whom received the Nobel Prize associated with these discoveries, **photovoltaic cells** that convert light into electricity, and the **GPS**.

Transistors are based on many quantum phenomena, particularly the electronic structure of atoms in semiconductors crystals that was discovered during the 1930s




Figure 53: How old were quantum scientists when they were awarded the Nobel prize in physics? © Olivier Ezratty, 2021.

and creates forbidden energy levels named band gaps (found by Sir Alan Herries Wilson, UK, in 1931), the impact of defects in crystals leading to doping and the tunneling effect due to the wave-particle duality of electrons. It also uses the field effect, which modulates the electrical conductivity of a material by the application of an external electric field. It was invented by Julius Edgar Lilienfeld (1882-1963, Austro-Hungarian and American) who filed in 1926 a patent using copper-sulfide semiconductor materials that was granted only in 1930, after a first filing in Canada in 1925. It corresponds to what we today call a “Field Effect Transistor” (FET). The first transistor invented in 1947 was made of germanium, not silicon. The first machine to run on transistors was the Manchester University Transistor Computer prototype, which operated in November 1953, followed by the Bell Labs TRADIC Phase One computer in 1954.

Transistors and lasers are the basis of much of today’s digital technology. All our digital devices are already quantum! The field of quantum optics started in the early 1960s with the laser invention and Roy J. Glauber’s work, with his seminal work in 1963 on light classification where he formalized the coherent states generated by lasers, *aka* Glauber states.

The post-war period was also dominated in quantum physics by advances made on superconductivity with the BCS theory in 1957 and the Josephson junction in 1962, and by the theoretical work of John Stewart Bell in 1964.

We then have the violation of Bell inequalities achieved by Alain Aspect’s experiments in 1982, which ruled out local hidden variables. 1980 and 1981 are other key

dates which mark the symbolic beginnings of quantum computing, imagined by Yuri Manin and Paul Benioff (gate-based quantum computing) and Richard Feynman (quantum simulation).

The term **second quantum revolution** covers advances from the 1990s and later, when the quantum properties of individual particles could be controlled at the level of photons (polarization, ...), electrons (spin) and atoms or ions, as well as superposition and entanglement. This led to the emergence of quantum cryptography and quantum telecommunications, in addition to the premises of quantum computing.



Felix Bloch (1905-1983, Swiss then American) is a physicist who created the geometrical representation of a qubit state in a sphere, Bloch’s sphere was elaborated in 1946 in a paper on nuclear magnetism, his main specialty.

Like other physicists of his time, he contributed to the Manhattan project, although quite shortly. He was awarded the Nobel prize in Physics in 1952 for his work on nuclear magnetic resonance and magnons conceptualization. He was also the first director of the international particle physics laboratory CERN in 1954.



Chien-Shiung Wu (1912-1997, Chinese then American) was a scientist who contributed to the development of nuclear physics and to the Manhattan project, with her gaseous diffusion process used for separating uranium 238 from uranium 235.




Figure 54: Timeline of key events in quantum physics after World-War II. © Olivier Ezratty, 2020-2025.


She is above all known for the experiment run at the end of 1956 and published in 1957 which demonstrated parity violation in the weak interaction with cobalt-60 nuclei, a result that earned Tsung-Dao Lee and Chen-Ning Yang the 1957 Nobel prize in physics while she was left out [121]. She also contributed to the development of quantum physics by conducting the first experiment on the angular correlation of photon pairs and entanglement in 1949, well before Alain Aspect's experiment in 1982 [122].

This experiment was different and was based on the measurement of the angular correlation of gamma ray photons (with very high-frequency and high-energy) generated by the encounter of electrons and positrons.



Hugh Everett (1930-1982, American) is a physicist who created the formulation of relative states and a global wave function of the Universe integrating observations, observers and tools for observing quantum phenomena.

He met Niels Bohr with other physicists in Copenhagen in 1959 to present his theory. He was politely listened to, but his interlocutors said that he understood nothing about quantum physics. Everett was also a contributor to the connections between the theory of relativity and quantum physics, especially around quantum gravitation. He is credited with the hypothesis of multiple or multiverse worlds, or many-worlds interpretation, explaining quantum entanglement and nonlocality.



The many-world interpretation actually comes from **Bryce DeWitt** (1922-2004, American) who extended Everett's work in 1970 and 1973 [123, 124]. DeWitt also worked on the formulation of quantum gravity theories.



John Wheeler (1911-2008, American) supervised Hugh Everett's thesis. He was specialized in quantum gravitation and nuclear physics. He worked in the Manhattan project, on the first American H-bombs, and on very high-density nuclear matter found in neutron stars.

He popularized the term black hole in 1967. He imagined a delayed-choice experiment to decide when a quantum object decides to travel as a wave or as a particle, which was later implemented in 2006 by a French team led by Philippe Grangier [125].

He collaborated with Niels Bohr and among his PhD students were Richard Feynman and Wojciech Zurek!



Richard Feynman (1918-1988, American) is one of the fathers of quantum electrodynamics with seminal papers published starting in 1948, which earned him the Nobel prize in Physics in 1965.

He introduced his famous Feynman diagrams which represent particle interactions, where lines represent particles and vertices represent interactions (Figure 55). He also developed the path integral formulation of quantum mechanics, which is based on the principle of least action and summing over all possible paths a particle could take. Before that, he was one of the contributors to the Manhattan project where, among other things, he cal-

culated with Hans Bethe the yield of nuclear explosions. He is also at the origin of the quantum explanation of helium superfluidity at very low temperature in a series of papers published between 1953 and 1958.





Figure 55: A Feynman diagram, which is used to describe particle and electromagnetic fields interactions and creation/destruction.

Richard Feynman shared his 1965 Nobel Prize in Physics with Julian Schwinger and Sin-Itiro Tomonaga who independently developed renormalization, the mathematical technique that eliminated the problematic infinite quantities from quantum electrodynamics calculations, allowing for accurate and verifiable predictions. All this also in 1948! This book will mention quantum algorithms that have been created for the simulation of the Schwinger model, a theoretical framework in quantum field theory that describes 1+1 dimensional quantum electrodynamics (QED), dealing with electrons and photons in a spacetime with one spatial dimension and one time dimension.

He theorized in 1981 the possibility of creating quantum simulators, capable of simulating quantum phenomena, which would be useful to design new materials and molecules in various fields like chemistry and biotech [126].

In 1985 [127], he described a model of gate-based quantum computing and how it could be equivalent to a quantum simulator solving a given Hamiltonian [128]. Richard Feynman was also famous for his great presentation skills and his many courses lecture notes.



Roy J. Glauber (1925-2018, USA) was a theoretical physicist, teaching at Harvard and at the University of Arizona. He got the Nobel prize in Physics in 2005 for his foundational work on the quantum theory of optical coherence.

He is considered to be a pioneer of non-classical light description and of the quantum optics field, with his work in 1963, describing various types of light (coherent, not coherent, ...). He also worked in the field of high-energy particle physics, which we don't cover in this book since out of scope of the "second quantum revolution".



Philip W. Anderson (1923-2020, USA) was a theoretical physicist who contributed to the theories of localization (*aka* "Anderson localization").

It describes extended states that can be localized by the presence of disorder in a system, antiferromagnetism and quantum spin liquid, symmetry breaking leading to the creation of the Standard Model, superconductivity (at high-temperature, pseudospin approach to the BCS theory, Anderson's theorem on impurity scattering in superconductors). He created the "condensed matter physics" naming. He got the Nobel prize in physics in 1977 for his work on the electronic structure of magnetic and disordered systems. He worked at the Bell Labs and was also a teacher at Cambridge University, UK.



John Stewart Bell (1928-1990, Irish) relaunched research in quantum mechanics in the 1960s on the notion of entanglement.

We owe him the Bell inequalities tests that highlight the paradoxes raised by quantum entanglement. Bell's 1964 theorem indicates that no theory of local hidden variables - imagined by Einstein in 1935 - can reproduce the phenomena of quantum mechanics [129]. In 1964, Bell was affiliated with the University of Wisconsin and worked at CERN in Geneva.

He was rather pro-Einsteinian in his approach and favorable to a realistic interpretation of quantum physics [130].

His Bell inequalities define the means to verify or invalidate the hypothesis of the existence of hidden variables explaining quantum entanglement. Bell's inequalities were violated by the experiments of **Alain Aspect** in 1982, demonstrating the inexistence of local hidden variables compatible with the formalism of quantum mechanics.

Prior to this experiment, Bell's inequalities had been formulated for pairs of entangled photons by **John Clauser** (1942, American, 2022 Nobel prize in physics), **Michael Horne** (1943-2019, American), **Abner Shimony** (1928-2015, American) and **Richard Holt** in 1969 with their so-called CHSH inequalities with some experimental settings proposals [131].

John Bell's work was completed in 2003 by **Anthony Leggett** (1938-2026, Anglo-American, Nobel prize in physics in 2003 for his work on superfluid helium) with his inequalities applicable to hypothetical non-local hidden variables [132]. Anthony Leggett was also an initial key contributor to what led to the creation of superconducting qubits. Anton Zeilinger (1945, Austrian) managed to experimentally violate these inequalities in 2007. According to Alain Aspect, however, this did not call into question the non-local hidden variable model proposed by David Bohm.

John S. Bell also created another theorem, in 1976, which refined the one from 1964. It demonstrated the incompatibility of quantum phenomena with the principle of local causality. In a nutshell, that quantum entanglement violates the principle of locality [133–135]. Entangled particles exhibit measurement correlations that seem to be instantaneous, even when separated by vast distances. Although it doesn't mean that some information is transmitted faster than light from one point to the other. This is still about some random measurement output. And when entanglement is used with the teleportation protocol, classical information links are used, which limits the speed of teleportation to the speed of this classical communication line which cannot be faster than light speed.



Klaus von Klitzing (1943, Germany) is a physicist known for discovery of the integer quantum Hall effect (QHE) in 1980. It led to a redefinition of the Ohm based on the Planck constant and the charge of the electron.

He made this discovery as an experimentalist while he was working at the Grenoble High Magnetic Field Laboratory in France (now the National Laboratory for Intense Magnetic Fields). He was awarded the 1985 Nobel Prize in Physics, and alone, which is rare nowadays. He created the von Klitzing constant, $R_K = h/e^2 = 25812.80745\dots$, e being the charge of the electron and h , the Planck constant.

The Hall effect, discovered by Edwin H. Hall in 1879 [136], describes the creation of a potential difference across a metal strip when an electric current flows through it and a perpendicular magnetic field is applied, causing electrons to deflect. While well-understood in common materials, the Hall effect exhibits novel behaviours in two-dimensional electron systems, which are artificially created thin layers of electrons. In these 2D systems, the Hall conductance doesn't change smoothly but instead shifts in discrete steps as the magnetic field is modified. These plateaus occur at conductance values that are integer multiples of e^2/h , in other words at Hall resistances equal to R_K divided by an integer.

The QHE (quantum Hall effect) discovery helped redefine the resistance standard independently of some reference material. The von Klitzing constant R_K is now precisely defined and used worldwide to calibrate resistance measuring instruments. It also led to the development of topological quantum computing, based on one of the many variations of the QHE (QSHE, QVHE, QAHE, FQAHE), the fractional quantum Hall effect (FQHE) that exhibits exotic states of matter with non-Abelian anyons. These quasiparticles are neither bosons nor fermions. Exchanging two of them applies a non-commuting operation on a degenerate ground state. Non-Abelian braiding statistics make them candidates for building topological qubits, which are resistant to decoherence and certain types of errors, a key requirement for fault-tolerant quantum computers. This is still

development in progress with very low TRL (technology readiness level).

The QHE helps in the characterization of the electronic properties of two-dimensional materials like graphene. The QHE was also used for the development of high-precision magnetic fields quantum sensors.

2.4 Quantum technologies physicists

This story now provides an overview of key contributors to the physics of quantum computing. They are often specialized in condensed matter, such as for superconducting qubits, and in photonics.

I highlight many European and French physicists, particularly those I have had the opportunity to meet for the last three years in my journey in the quantum ecosystem. This inventory is both objective and subjective. Objective because it includes a broad and worldwide hall of fame in the field. Subjective because I have added a good dose of physicists I know. It creates a measurement bias which is easy to understand in social science as well as in quantum physics.

Starting in the 2023 edition, I have clustered these physicists in broad categories, the first ones being the generalists who worked in different fields. See another simplified time of the recent period in Figure 56.

2.4.1 Generalists



Alain Aspect (1947, French, 2022 laureate of the Nobel prize in physics) observed the violation of Bell's inequalities with a series of experiments conducted between 1980 and 1982 at the Institut d'Optique with Jean Dalibard, Philippe Grangier and Gérard Roger.

Taking the principles of quantum physics for granted, it validated the nonlocality of quantum properties and ruled out broad classes of local hidden-variable theories [137].

Alain Aspect's experiments were using calcium atoms as source of photons, using some laser excitement and an atomic cascade generating pairs of entangled photons in the visible spectrum at 551 nm and 423 nm. There were actually several experiments: in 1981 with Philippe Grangier and Gérard Roger with one way polarizers, 1982 also with Grangier and Roger with two-channels polarizers and also 1982, with Jean Dalibard and Gérard Roger, using variable polarizers based on acousto-optical 10 ns switches. These could act faster than light propagation between the polarizers (40 ns) and even than the photons time of flight between the source and each switch (20 ns). It avoided any potential synchronization between the polarizers, using acousto-optical switches driven at around 50 MHz on both sides, feeding two orthogonal polarizers and photon detectors, using different




Figure 56: Quantum computing key events timeline from 1990 to 2020. Many other events happened since then which are described in this book. In green, experimentalists and experiments. © Olivier Ezratty, 2020-September 2026.

angles (Figure 57). Because the switching pattern was periodic rather than generated independently at random, a setting-independence loophole remained. Later experiments strengthened spacelike separation and closed detection loopholes [138, 139]. Aspect published a book in 2025 describing the whole story on these experiments and the many left open questions it created [140].

From 1988 to 2015, other experiments were conducted elsewhere and implemented loophole-free Bell tests, first closing individual loopholes and then, in 2015, closing them altogether. It confirmed then that no local variable theory could explain the observed correlations and validated one key nonlocality condition: having a long distance between the photon analyzers to avoid any interactions made possible by special relativity. It avoided potential detection loopholes with high-efficiency photon detectors on top of escaping ‘memory loopholes’, which was already obtained by Alain Aspect et al. in their seminal 1982 experiment [141, 142]. A loophole-free Bell inequality violation experiment was even achieved with superconducting circuits in 2023 [143].

After his work on photon entanglement, Alain Aspect shifted gear on cold atoms control with lasers, starting with helium. This led to the creation of a promising field of quantum computing in France, using cold atoms, embodied by the startup **Pasqal**, whose scientific director is Antoine Browaeys, a former PhD student of Alain Aspect who also worked with Philippe Grangier. Along with other scientists, Alain Aspect is also a member of Eviden/Atos Scientific Council and in the scientific board of **Quandela**.

He used to teach quantum physics, notably in MOOCs created for Ecole Polytechnique and distributed by Coursera. One of his prides is to see his former PhDs and post-docs becoming entrepreneurs and turning sci-

ence into useful products within the framework of the second quantum revolution.



Philippe Grangier (1957, French) was the first PhD student of Alain Aspect with whom he worked on the series of Aspect’s experiments with Gérard Roger and Jean Dalibard.

In 1986, he demonstrated nonclassical anticorrelation at a beam splitter [144]. He created the foundations of CV-QKD with Frédéric Grosshans in 2002. He was then involved in the creation of the associated startup, Se-QuireNet, in 2008 and closed in 2017, probably created a little too early in relation to the needs of the market. He was invested in cold atoms control with lasers at IOGS (Institut d’Optique). He cocreated the CSM ontology of quantum foundations with Alexia Auffèves and Nayla Farouki, starting in 2013 and with a series of foundational papers published between 2015 and 2019. CSM ontology is quickly covered in the quantum foundations section.

He was also active in the French and European quantum ecosystems, with driving the DIM SIRTEQ in the Paris Region from 2017 to 2020 (a regional fund for academic research in quantum science), and also ran the European Quantum Flagship coordination project QUCATS until July 2025.



Jean Dalibard (1958, French) is a research physicist at the ENS and teacher at Ecole Polytechnique and the Collège de France. He is a specialist in quantum optics and interactions between photons and matter.

He participated with Philippe Grangier in the set-up of Alain Aspect’s experiment in 1982 when he was a contingent scientist at the Institut d’Optique. He cre-




Figure 57: Alain Aspect et al. 1982 Bell inequality test experiment setup. Comments, Olivier Ezratty, 2021-2025.

ated the idea of the magneto-optical trap (MOT) in 1987 that is used to cool neutral atoms using a mix of variable magnetic fields and lasers [145].

In 1989 with Claude Cohen-Tannoudji, he created the “Sisyphus cooling” technique, a polarization gradient cooling technique which cools atoms below the Doppler cooling limit, down to the recoil limit.



Anton Zeilinger (1945, Austrian, 2022 Nobel prize in physics) is a physicist who advanced the field of quantum teleportation in the 2000s. He also demonstrated in 1988 the diffraction of neutrons through single and double slits.

His team was the first to demonstrate photon quantum teleportation, in Innsbruck in 1997 [146]. He is a specialist in quantum entanglement, having proved that it is possible to entangle more than two quantum objects or qubits. He created theoretical and experimental foundations for quantum cryptography.

With two colleagues, he also developed the GHZ (Greenberger-Horne-Zeilinger) enables yet another demonstration of the inexistence of hidden variables which would explain quantum entanglement of at least three particles and with a finite number of measurements.

The concept was created in 1989 and was validated experimentally in 1999. Anton Zeilinger also supervised the thesis of **Jian-Wei Pan**, who became later the

quantum research czar in China with the development of many advances, particularly in quantum communications and photonics.



Frank Wilczek (1951, American) is a professor of physics at MIT and the chief scientist at the Wilczek Quantum Center in Shanghai. He was awarded in 2004 the Nobel Prize in Physics, shared with David Gross and H. David Politzer, for his work on the theory of strong interaction and quantum chromodynamics.

He is known for his work on quasiparticles and anyons in 1982. He also proposed the notion of continuous time crystals in 2012 [147]. One of his former PhD students is Chetan Nayak who runs the Microsoft hardware team developing Majorana fermion-based qubits (as of mid-2026).



John Preskill (1953, American) is a professor at Caltech. On top of countless scientific contributions, he is the creator of the quantum supremacy notion, introduced in a talk at the 2011 Solvay conference and in the resulting 2012 paper.

He also created the NISQ expression in 2018, for Noisy Intermediate-Scale Quantum, related to noisy quantum computers. He is a regular speaker at conferences where he reviews the state of the art of quantum computing.

He is also involved with Amazon and their cat-qubits superconducting project revealed in December 2020.

He cofounded the startup Oratomic in March 2026 with Dolev Bluvstein (CEO), Manuel Endres and Hsin-Yuan Huang (CTO) among others, a Pasadena based neutral atom venture which raised a \$300M series A in July 2026. The company had published a resource estimate which would break RSA-2048 keys with Shor integer algorithms with 10,000-26,000 atoms.



Jian-Wei Pan (1970, China) is the leading quantum physics scientist in China. He is a professor and Executive VP at USTC (University of Science and Technology of China) and a member of CAS (China Academy of Science).

He did his PhD in Vienna under the supervision of Anton Zeilinger. He and his team are famous for premiere experiments on photons quantum entanglement in 2004, quantum key distribution over a satellite (2017), with boson sampling (Jiuzhang, December 2020) and superconducting qubits (Zuchongzhi, 2021). He plays a key role coordinating China's quantum efforts at CAS.



Dieter Zeh (1932-2018, German) is the discoverer of the quantum decoherence phenomenon in 1970. It describes the progressive end of the phenomenon of superposition of quantum states, when a system becomes entangled with uncontrolled environmental degrees of freedom.

The notion of decoherence is key in the design of quantum computers. The objective is to delay it as much as possible resulting from the interaction between quanta and their environment [148].



Wojciech Zurek (1951, Polish) is a quantum decoherence physicist who contributed to the foundations of quantum physics applied to quantum computers.

We owe him the no-cloning theorem, published in 1982 with William Wootters and, independently, by Dennis Dieks, which states that no unitary operation can produce an exact copy of an arbitrary unknown quantum state [149, 150]. He is also at the origin of the concept of quantum Darwinism which would explain the link between the quantum world and the macrophysical world.



Maciej Lewenstein (1955, Polish) is a theoretical physicist, specialized in quantum optics of dielectric media and cavity quantum electrodynamics, teaching at ICFO in Spain.

He worked with many leading worldwide scientists including Roy J. Glauber (Nobel in Physics in 2005) at Harvard, Thomas W. Mossberg, Andrzej Nowak, Bibb Latané, Anne L'Huillier (CEA, France), Peter Zoller and Eric Allin Cornell (Nobel in Physics in 2001 for

his work on Bose-Einstein condensates in 1995), in the USA, France, Spain, Poland and Germany.

His contributions span an incredible number of fields like the physics of ultra-cold gases, quantum information, quantum optical systems, quantum communications, quantum cryptography, quantum computers, mathematical foundations of quantum physics, tensor networks and entanglement theory, laser-matter interactions attosecond physics, quantum optics (cQED), atoms cooling and trapping, non-classical states of light and matter and quantum physics foundations.



David Mermin (1935, American) is a solid-state physicist at Cornell University known for the Hohenberg–Mermin–Wagner theorem and for his use of the term “boojum” to describe the phase patterns of superfluidity.

He also worked on quantum physics foundations on Quantum Bayesianism, or QBism, and quantum information science [151]. He coined the expression “*shut up and calculate*” as a simple way to describe Niels Bohr Copenhagen interpretation of quantum physics [152].

2.4.2 Cold atoms



Claude Cohen-Tannoudji (1933, French) is a former student of Ecole Normale Supérieure (ENS Paris) where he followed the teachings of mathematicians Henri Cartan and Laurent Schwartz and physicist Alfred Kastler.

He was awarded the Nobel prize in Physics in 1997, along with Steven Chu (who became the Secretary of Energy during Barack Obama's first term), and Bill Phillips. Bill Phillips was awarded the Nobel prize in physics for his work on the Zeeman slower in 1982 which later enabled Bose-Einstein condensation experiments.

Claude Cohen-Tannoudji owes his Nobel Prize to his work on atoms cooling with lasers which made it possible to reach extremely low temperatures, below the milli-Kelvin. Alain Aspect once worked in his team at Ecole Normale Supérieure in Paris.

Alain Aspect says that he discovered quantum physics with reading the reference book on quantum physics by Claude Cohen-Tannoudji, Bernard Diu and Franck Laloe published in 1973 [153–155]. It is split into three volumes that were updated in 2019 and has a total of over 2,300 pages. These are classical quantum physics student textbooks throughout the world.



Bill Phillips (1948, American) was awarded the Nobel prize in physics in 1997 for his work on the Zeeman slower which enabled Bose-Einstein condensation experiment.

Nowadays, it is the first process used for cooling cold atoms in a 1D cylinder using a mix of counterpropagating laser excitation implementing Doppler's effect, and a spatially inhomogeneous variable magnetic field, before injecting them in a magneto-optical-trap (MOT) which confines the atoms in a narrow location.



Serge Haroche (1944, French), Nobel prize in Physics in 2012, is a founder of Cavity Electrodynamics (CQED) which describes the interaction between photons and atoms in cavities. He used it to create cold atom based qubits.

Jean-Michel Raimond and **Michel Brune** were among his key collaborators. Serge Haroche was the first to measure the phenomenon of quantum decoherence (loss of superposition) of a mesoscopic superposition of cavity-field states in a 1996 experiment. This experiment was using rubidium atoms. Serge Haroche is also a member of Atos Scientific Council.

CQED was later applied in the field of superconducting qubits with Circuit Electrodynamics (cQED), where atoms are replaced by an artificial atom made with a Josephson junction and the cavity by a planar microwave resonator. Serge Haroche is one of the most circumspect scientists on the future of quantum computing, at least for universal gate computing. He believes more in the advent of quantum simulation [156].

Other scientists brought key contributions in atoms science. **Daniel Kleppner** (1932, American) achieved in 1998 the first Bose-Einstein condensate with atomic hydrogen, the species originally expected to condense first, three years after Eric Cornell and Carl Wieman got there with rubidium at JILA [157]. **Herbert Walther** (1935-2006, German) did pioneering work in cavity quantum electrodynamics and also with trapped ions. He created the Max Planck Institute of Quantum Optics in 1981. **Gerhard Rempe** (1956, German) developed cavity quantum electrodynamics with the control of neutral atoms using microwaves, in connection with **Jeff Kimble** (1949, American, Caltech).



Christophe Salomon (1953, French) is a physicist specialized in photonics and cold atoms, research director at the LKB (ENS Paris). He is particularly interested in quantum gases superfluidity (Bose-Einstein condensates) and in time measurement with cesium atomic clocks.

He did a thesis in laser spectroscopy and then did a post-doc at the joint JILA laboratory between NIST and the University of Colorado. He is also a member of the Academy of Sciences since 2017.



Immanuel Bloch (1972, Germany) is the scientific director at the Max Planck Institute of Quantum Optics, Garching and professor for experimental physics at the Ludwig-Maximilians University (LMU) in Munich.

He is a leading experimentalist in the cold atom domain, his work covering quantum gases in optical lattices, the first realization of a quantum phase transition from a weakly interacting superfluid to a strongly interacting Mott insulating state of matter based on a theoretical proposal by Peter Zoller and Ignacio Cirac. He also works on fermionic simulation using analog cold-atom experiments.



Mikhail Lukin (1971, Russian and then American) is a quantum physics professor at Harvard. He's a prolific scientist with a skyrocketing h-index of 213 (as of September 2026).

He is working on quantum optics, quantum control of atomic and nanoscale solid-state systems, quantum sensing, nanophotonics and quantum information science. He is behind many feats in cold atoms physics with record two-qubit gate fidelities of 99.5% obtained in 2023 with 60 atoms as well as in the NV centers field, being the co-inventor in 2008 of NV centers based magnetometry [158, 159].

He cofounded QuEra (USA) that develops a cold atoms gate-based quantum computer, reaching 256 qubits as of 2025 (in analog mode), and created many new advances since 2023 for the development of gate-based cold-atom computing. He is also a cofounder and scientific advisor of QDTI (USA).



Marie-Anne Bouchiat (1934, French) is a specialist in rubidium atoms physics and their control by optical pumping. This is the basis for the creation of quantum computers based on cold atoms.

Her daughter **Hélène Bouchiat** (1958, French) is also a physicist, specialized in condensed matter at the LPS laboratory of the University Paris-Saclay and member of the Académie des Sciences since 2010, like her mother who has been there since 1988.



Antoine Browaeys (1972, French) is a CNRS research director leading the quantum optics-atom team in the Charles Fabry Laboratory at Institut d'Optique specialized in the control of cold atoms.

He is a pioneer in the control of individual cold atoms and on their use in quantum simulators and digital quantum computers. He is also a cofounder and the scientific director of Pasqal, a startup designing a cold atoms computer used as a quantum simulator, and then, as a gate-based quantum computer. He was awarded the CNRS silver medal in 2021 and the 9th Bell prize in 2025.



Elisabeth Giacobino (1946, French) is a specialist in laser physics, nonlinear optics, quantum optics and superfluidity, particularly in relation to the control of cold atoms.

She worked at the CNRS in the ENS LKB (Laboratoire Kastler-Brossel). She was a member of the scientific selection committee of the European Quantum Flagship and also for the ANR (Agence Nationale de la Recherche).



Hélène Perrin (1972, French) is CNRS research director working at the Laboratoire de Physique des Lasers (LPL) from Université Sorbonne Paris Nord, working on Bose-Einstein condensates and cold atoms control.

Together with Pascal Simon, she drives the Quantum Simulation SIM project, a cold atom-based quantum simulator. She also gives lessons on quantum computing. She did her PhD thesis with Christophe Salomon at the ENS LKB in Claude Cohen-Tannoudji's group. At CEA-Saclay, she also worked on fractional quantum Hall effect. Since 2022, she is the director of QuanTIP, the Paris region quantum ecosystem network.



Francesca Ferlino (1977, Italian) is a typically European researcher, having worked in many laboratories from different countries. She is research director at the IQOQI in Innsbruck, Austria, where she leads the Dipolar Quantum Gases laboratory.

She is a specialist in cold atoms and erbium-based Bose-Einstein condensates.

2.4.3 Trapped ions



Wolfgang Paul (1913-1993, German), not to be confused with Wolfgang Pauli, is a physicist who conceptualized trapped ions in the 1950s. He got the Nobel prize in physics in 1989. We owe him the traps that bear his name and are used to control trapped ions.

He shared his Nobel prize with **Hans Georg Dehmelt** (1922-2017, German then American), who developed the complementary Penning trap using a static magnetic field, and with **Norman Ramsey** (1915-2011, American) for the separated oscillatory fields method. In Innsbruck, Austria, the physicists **Juan Ignacio Cirac** (1965, Spanish) and **Peter Zoller** proposed in 1995 the first realistic scheme for a trapped ion quantum computer, based on the work of Wolfgang Paul [160]. **Rainer Blatt** and his team implemented that specific Cirac-Zoller CNOT gate in 2003.

The very first quantum logic gate had been demonstrated in the meantime, at the end of 1995 at NIST Boulder by Christopher Monroe and David Wineland, on a single beryllium ion whose motional and internal states played the roles of the control and target qubits, just a year after Peter Shor created his famous eponymous integer factoring algorithm [161].



David Wineland (1944, American) is a Boulder-based NIST physicist known for his advances in trapped ions and their laser-based cooling in 1978.

He also created in 1995 the first single quantum gate operating on a single atom. He was awarded the Nobel prize in Physics in 2012 jointly with Serge Haroche for his advances in atoms and ions laser cooling, a technique he first experimented in 1978, followed by the first quantum gate applied to a trapped ion in 1995 and the entanglement of four trapped ions in 2000.



Rainer Blatt (1952, Austrian and German) from the University of Innsbruck is an experimental physicist specialized, among other things, in trapped ions qubits. His team created an eight-ion W state in 2005 and a fourteen-ion GHZ state in 2011, deterministic two-ion entanglement having been obtained first at NIST Boulder in 1998.

He co-founded Alpine Quantum Technologies (AQT) with Thomas Monz (CEO). The company's ambition is to create and commercialize a trapped ions based quantum computer. He also works at TUM in Munich, Germany and contributed to the launch of the Munich Quantum Valley starting in 2021.



Christopher Monroe (1965, American) is an American physicist known for his work on trapped ions and for co-founding IonQ in 2015, one of the two best funded quantum startups worldwide with PsiQuantum.

He worked on trapped ions with David Wineland at the NIST laboratory in Boulder, Colorado, before setting up his own group at the University of Maryland. He demonstrated the ability to trap ions, create ion-based quantum memory and create analog quantum simulators. He also ran a laboratory at the University of Michigan in the early 2000s. He now works at Duke University.



Tracy Northup (1978, Austrian) is a researcher working on trapped ions and optical cavities, one of the major branches of quantum computing.

She leads the Quantum Interfaces Group laboratory at the University of Innsbruck, which is one of the most active in the field of trapped ions, a major Austrian specialty. Among other things, she developed fiber based cavities to trap ions and enable ions-photon coupling [162].



Perola Milman (c. 1975, French) is a specialist in the theory of quantum computing and in particular with trapped photons and ions. In particular, she has demonstrated capacity to entangle molecules.

She is a lecturer-researcher at the Laboratory of Quantum Materials and Phenomena of the University Paris Diderot and a professor on the quantum theory of light and quantum entanglement.

2.4.4 Superconductivity



Brian Josephson (1940, English) is a physicist from the University of Cambridge. He was awarded the Nobel prize in Physics in 1973 at the age of 33, for his prediction in 1962 of the effect that bears his name when he was only 22 years old and a PhD student at the University of Cambridge.

Brian Josephson shared the 1973 Nobel prize in physics with two scientists who had worked before him in the same field: Leo Esaki (1925, Japan) for his discovery of the tunnel effect in semiconductors in 1958 and Ivar Giaever (1929, Norway, also still alive) who found that this effect could occur in superconducting materials in 1960. The Josephson effect describes the passage of current in a superconducting circuit through a thin insulating barrier a few nanometers thick, using tunneling effect, and the associated threshold effects. This effect has a continuous and alternative current variant, depending on the voltage applied to the junction (Figure 58).

It is generated by electrons with opposite spins organized in Cooper pairs named after Leon Cooper who described them in 1956. These pairs behave as bosons. These electrons pairs have opposite spins (magnetic polarity). The system behaves as a resistance associated with a loop inductance, the oscillation being controllable by a magnetic field and having two distinct energy states and phase. Superconductivity was discovered in 1911 by **Heike Kamerlingh Onnes** (1853-1926, the Netherlands). This is the basis of superconducting qubits and their quantum gates!




Figure 58: The alternative current Josephson effect is based on multiple Cooper pairs of opposite spin electrons that are traversing a potential barrier, using an oscillating current. The usual barrier is made of aluminum oxide while the superconducting material is made of aluminum. © Olivier Ezratty, 2021.



Daniel Esteve (1954, French) is a physicist in charge of the CEA's Qnantronics laboratory in Saclay, France, launched in 1984 with Michel Devoret and Cristian Urbina, and part of the IRAMIS laboratory. He contributed to the development of transmon superconducting qubits.

With Vincent Bouchiat, he demonstrated in 1998 a first Cooper pair box showing quantum coherence, followed in 2002 by the qnantrium designed with Denis Vion, a charge qubit operated at an optimal working point where it is first-order insensitive to charge and flux noise [163].



Michel Devoret (1953, French) is a telecom engineer turned physicist, cofounder of the Qnantronics laboratory with Daniel Esteve at the CEA in Saclay, which he co-led from its 1984 creation until 1995.

He was a professor at Yale University from 2002 to 2023 and a cofounder of QCI with his colleague Rob Schoelkopf (1964, USA), which he left in 2019/2020. He worked several times with John Martinis, when John was a PhD student in UCSB with John Clarke, then when he was a post-doc at CEA in Saclay, and at last at UCSB, where they wrote together a review paper in 2004 on superconducting qubits [164]. In 2023, Devoret moved to UCSB and started working with Google, contributing to the development of the Willow chip and reducing correlated errors with the use of gap engineering of the qubit Josephson junction.

John Clarke, John Martinis and Michel Devoret were awarded the Nobel prize in physics in October 2025.



Steven Girvin (1950, American) is a professor of physics at Yale University, specialized in condensed matter physics, and Director of the Co-design center for Quantum Advantage, at Brookhaven National Laboratory since 2020.

He is a key contributor to works on circuit quantum electrodynamics (cQED) and superconducting qubits. At Yale, he works with Robert Schoelkopf and Michel Devoret on the various engineering problems associated with superconducting qubits.



Rob Schoelkopf (1964, USA) a physicist and director of the Yale Quantum Institute. Along with Steve Girvin and Michel Devoret, he made key advances in superconducting qubits. He particularly worked on single-electron devices, being the inventor of the Radio-Frequency Single-Electron Transistor.

He also created the field of circuit quantum electrodynamics (cQED) with Andreas Wallraff and Alexandre Blais who were respectively Yale post-doc and PhD student around 2002-2004. In 2007, with Steven Girvin, he engineered a superconducting communication bus to store and transfer information between distant qubits

on a chip. In 2009, their team, also including Alexandre Blais and Jay Gambetta, demonstrated the quantum processor running some quantum computation, with two qubits [165]. At last, he cofounded Quantum Circuits Inc in 2015, a startup developing superconducting qubits using a dual rail technology, with a better protection against errors.



John Martinis (1958, American), is a physicist from UCSB who famously worked at Google between 2014 and 2020 where he led the hardware team in charge of superconducting qubits up to creating the Sycamore processor and its related “quantum supremacy experiment” in 2019.

After his thesis at Berkeley with John Clarke on what later became superconducting qubits, he did a post-doc in Daniel Esteve’s Qnantronics laboratory at the CEA in Saclay. He cofounded Qolab (2024, USA), to develop utility scale superconducting qubits quantum computers. He was awarded the Nobel prize in physics in 2025.



Jay Gambetta (1979, American) is the scientist leading as a VP since 2019 IBM’s research team working on superconducting qubits quantum computers after running the IBM team that created and launched IBM Quantum Experience, Qiskit and the IBM Quantum System One in 2019.

He joined IBM in 2011. After a thesis in quantum foundations and non-Markovian open quantum systems done in Australia in 2004, he focused on developing superconducting qubits, first in a post-doc tenure at Yale University and then at the Institute for Quantum Computing in Waterloo. He also worked on quantum validation techniques, quantum codes and applications.



Andreas Wallraff (German) is a Professor for Solid State Physics at ETH Zurich after having obtained degrees in physics from the London Imperial College and RWTH Aachen in Germany and worked at the Jülich Research Center also in Germany, Yale University in the USA and the LKB in France.

He is specialized in the coherent interaction of single photons with quantum electronic circuits and quantum effects as well as on hybrid quantum systems combining microwave control, superconducting circuits and semiconductor quantum dots.



Alexandre Blais (Canada) is a Professor in the Department of Physics and Director of the Université de Sherbrooke’s Institut Quantique. He is one of the key contributors to the development of circuit quantum electrodynamics (cQED) that enable the creation of superconducting qubits.

He is also a cofounder of Nord Quantique, a Quebec startup developing bosonic code qubits. Like Jay Gam-

betta, he did a post-doc at Yale, the US epicenter of the early developments of superconducting qubits.



Irfan Siddiqi (1976, American Pakistani) is one key contributor to advancements in superconducting qubits. He did his PhD and post-doc at Yale, working initially in aluminum hot-electron bolometers for microwave astronomy and then, high frequency measurement techniques for superconducting qubits.

He developed the Josephson Bifurcation Amplifier that uses the non-dissipative and nonlinear nature of the Josephson junction to create high gain and minimal back action readout of qubits. This led to the creation of superconducting parametric amplifiers and Josephson traveling wave parametric amplifiers. He then moved at Berkeley University and the DoE Lawrence Berkeley National Laboratory. He works on quantum electrodynamics, quantum error correction, multi-partite entanglement generation and single photon detection. He runs there the Advanced Quantum Testbed, an integrated research platform on superconducting qubits and enabling technologies.



Patrice Bertet (c. 1976, France) is part of Daniel Esteve’s team at CEA-SPEC. He did his thesis at Serge Haroche on Rydberg atoms and then went to Delft University.

He participated in the early days of superconducting qubits (quantronium at CEA and TU Delft). He then worked on QED (quantum electrodynamics) circuits based on cavities and then on transmon qubits. He is working on the association of superconducting qubits and the measurement of their state with electron spins, notably based on NV centers, which can also be used for quantum memories.



Audrey Bienfait (c. 1990, France) is a former PhD student of Patrice Bertet at CEA-SPEC who is now doing her research at ENS Lyon in the team of **Benjamin Huard** (1979, French).

She was awarded the Bruker Prize 2018 for her thesis on electron paramagnetic resonance or “ESR - Electron Spin Resonance” in quantum regime and the Michelson Postdoctoral Prize 2019 in March 2020 for her work on the entanglement of superconducting qubits via phonons.



Shi Yaoyun (1976, Chinese) is a professor at the University of Michigan who led the Alibaba Quantum Laboratory and its fluxonium superconducting qubit effort until Alibaba closed the lab in November 2023 and donated its equipment to Zhejiang University.

He created various records of quantum simulation on server clusters that we will describe in this book. He earned a computer science PhD from Stanford. He

also worked on quantum cryptography and certifiable randomness.

2.4.5 Spin qubits



Daniel Loss (1958, Swiss) proposed in 1998 with David DiVincenzo to use electron spins in quantum dots to create a quantum computer. He currently is the Co-Director and founding member of National Center on Spin Qubits (NCCR SPIN).

It gathers the University of Basel, EPFL and IBM Zurich, an initiative from the Swiss Nanoscale Center SNI. He is a professor at the University of Basel as well as at the King Fahd University of Petroleum and Minerals in Saudi Arabia. After a PhD in theoretical physics at the University of Zurich in 1985 he was a post-doc in the group of Anthony J. Leggett in the USA and at IBM Research. After a stint in Vancouver, he went back to Switzerland. He works on condensed matter physics and spin-dependent and phase-coherent phenomena in semiconducting nanostructures and molecular magnets with applications in quantum computing. In 2026, he proposed an energy efficient spin-based architecture to implement semi-classical reversible computing [166].



Bruce Kane (c. 1958, American) is a researcher at the Joint Quantum Institute from the University of Maryland (a JV with NIST).

While he was doing research at UNSW, he presented in 1998 the atom “donor spin” model, a spin-based qubit concept based on using individual phosphorous atoms in pure silicon lattice structures. This is the principle on which Michelle Simmons works at both UNSW and her startup SQC. The jury’s still out to demonstrate that this technology can scale among the various spin qubits proposals.



Leo Kouwenhoven (1963, Dutch) is a quantum physicist who got his PhD at TU Delft in 1992 and became a professor there in 1999. He led experimental results on the potential “signatures” of Majorana fermion quasiparticles in 2012 and later on their “definitive” existence in 2018.

The related Nature paper was retracted in 2021 due to experimental data mismanagement and reporting. From 2016 till 2022, he was a researcher at Microsoft Research, which he left in 2022, returning to his home based at QuTech and the Kavli Institute of Nanoscience from TU Delft. He then continued to contribute to Majorana related research [167–169].



Lieven Vandersypen (1972, Belgian) started as a mechanical engineer and a PhD at Stanford, then went to IBM in Almaden, California, where he became interested in MEMS.

He demonstrated the use of Shor’s algorithm for factoring the number 15 with NMR qubits, and then became a researcher at TU Delft University in the Netherlands and in its QuTech spin-off, which he currently runs. He is a pioneer of electron spin qubits. In this capacity, he works notably with Intel, and is testing their FinFET-based qubit chips at QuTech with Intel, which invested \$50M in QuTech in 2015. He participated to the creation of the startup Groove Quantum in 2023, which develops silicon-germanium spin qubits quantum computers.



Menno Veldhorst (1984, Dutch) is a group leader at QuTech. He got his PhD in 2012 on superconducting and topological hybrids at the University of Twente. He then worked on silicon quantum dots at UNSW.

There, he demonstrated in 2015 the first two qubit operations in silicon. At QuTech, he works on silicon and silicon/germanium (SiGe) qubits to build scalable quantum computers. His team is currently pioneering work on SiGe/Ge qubits with qubits manipulation in arrays of up to 16 quantum dots. He proposed a crossbar array architecture to create logical qubits.



Michelle Simmons (1967, British-Australian) is a physicist from the University of New South Wales in Australia (UNSW), working on silicon spin qubits.

She is the director of CQC2T (Centre of Excellence for Quantum Computation and Communication Technology) from UNSW. She is also the cofounder of SQC (Silicon Quantum Computing), the leading quantum computing Australian startup (\$66M), a spin-off from her university and from CQC2T.

In 2019, her team built the first two-qubit gate between phosphorous atom qubits in silicon, operating in only 0.8 ns. It became a full-fledged 10 qubit processor in 2022. She is using STM (scanning tunneling microscopes) to position phosphorus dopants in the silicon substrate.



Andrew S. Dzurak (Australian) is the Director of the Nanotechnology Fabrication Unit at UNSW’s Australian National Fabrication Facility from the CQC2T research center.

This facility’s white room is used to manufacture silicon qubits chips. Andrew Dzurak has been a pioneer of silicon qubits since 1998. He is leading research at CQC2T on silicon qubit control and reading. He created the first phosphorus-based silicon double qubits in 2015. He was a lead scientist for SQC, founded by Michelle Simmons, but seemingly left the company in 2021.

He created Diraq in 2022, a startup dedicated to the creation of scalable quantum computers using quantum dot silicon spin qubits.



Andrea Morello (1972, Italian) is one of the star researchers at UNSW in Australia. He is Program Manager of the ARC Centre of Excellence at CQC2T and leads the Fundamental Quantum Technologies Laboratory at UNSW.

During his studies, he attended the Laboratoire National des Champs Magnétiques Intenses of the CNRS in Grenoble. Today he is one of the specialists in silicon-based qubits. He is also a quantum engineering teacher at UNSW. His team was the first to demonstrate coherent control and readout of an individual phosphorus atom electron and nuclear spin in silicon and held for many years the record for the longest quantum memory time of 35.6 s in a single solid-state qubit.



James Clarke (c. 1971, American) launched Intel's quantum computing research efforts and the Director of Quantum Hardware at Intel since 2015. He's also behind Intel's partnership with QuTech in The Netherlands.

He is currently focused with his team of about 100 researchers and engineers on creating scalable quantum computers with silicon and SiGe qubits. He started working at Intel as a process engineer in 2001 after having studied and worked on organic chemistry (PhD in Harvard and post-doc at ETH Zurich).



Maud Vinet (1975, French) started as physics engineer and was granted a PhD in physics from Grenoble University. She then spent 20 years working in silicon technologies development and transfer for the semiconducting industry.

She led the silicon qubit project at CEA-Leti in Grenoble. Since 2016, CEA-Leti was focused on silicon spin qubits leveraging the strong relationships between fundamental science and technology in Grenoble ecosystem. In November 2022, Maud Vinet launched Quobly with Tristan Meunier (CNRS) and François Perruchot (CEA-Leti). Maud was driving QLSI, the European Quantum Flagship research project on silicon spins qubits, awarded in March 2020, after obtaining with **Tristan Meunier** (1977, French, at CNRS Institut Néel) and **Silvano de Franceschi** (1970, Italian, at CEA IRIG) an ERC Synergy grant funding of €14M in 2018 for the QuCube silicon qubit project.

Beforehand, she had previously contributed to the industrialization of FD-SOI technology with CEA and STMicroelectronics, GlobalFoundries and IBM. FD-SOI (Fully-Depleted Silicon on Insulator) is using a layer of silicon oxide insulator and channels of undoped silicon between the drain and the source, limiting leakage between the latter two.



Anne Matsuura (c. 1970, Japanese American) is a physicist who is leading the Quantum & Molecular Technologies team from the Intel Quantum Research Laboratory since 2014.

She leads the American's efforts in the creation of superconducting and silicon qubits quantum computers, with an overall vision of the hardware architecture. Her impressive career starts with a thesis at Stanford in synchrotrons, then in US Air Force labs and In-Q-Tel (the CIA investment fund). She also directed the European Theoretical Spectroscopy Facility in Belgium.


Let's also mention **Matthieu Desjardins** (C12), **Takis Kontos** and **Audrey Cottet** (LPENS) who are behind the development of carbon nanotubes qubits at C12 in France.

2.4.6 NV centers



Jean-François Roch (1964, French) is a quantum physics professor at ENS Paris Saclay. He is a pioneer of the usage of NV centers in many applications.

Particularly in quantum sensing, including for studying matter and magnetism at very high-pressure, which could be helpful for the discovery of high-temperature superconducting materials. He conducts this research in partnership with Thales and with the CEA. He also led the founding Wheeler delayed choice experiment in 2006.



Ronald Walsworth (c. 1972, American) is a pioneer in the usage of NV centers for quantum sensing in various fields, from life sciences to physics and astrophysics like for the detection of dark matter.

He leads the Walsworth group at the University of Maryland and is the founding director of the UMD Quantum Technology Center. Several startups emerged from his lab like qdm.io, Hyperfine.io (MRI) and QDTI (which he both cofounded).

He also launched the Quantum Catalyzer quantum startups accelerator (Q-CAT) that creates quantum startups from scratch. He got a PhD in physics from Harvard in 1991.



Ronald Hanson (1976, Dutch) is an experimental physicist from TU Delft known for his work on the foundations and applications of quantum entanglement, and his work with NV centers. He is also the scientific director of QuTech.

2.4.7 Photonics



Jürgen Mlynek (1951, German) is a physicist specialized in optronics and interferometry. He was the coordinator of the strategic advisory board behind the launch of the European Flagship project on quantum in 2018.

We owe him, as mentioned in connection with Louis De Broglie, the experiment validating the wave-particle duality of atoms carried out using helium in 1991 with Olivier Carnal at the University of Konstanz.



Jean-Michel Gérard (1962, French) is a physicist from the CEA IRIG laboratory in Grenoble and director of the joint PHELIQS laboratory (PHotonics, ELectronics and Quantum Engineering) from UGA (University of Grenoble) and CEA.

He works in particular on the creation of single photon sources based on quantum dots as well as single photon detectors based on superconducting nanowires and OPO laser diodes.



Andrew G. White (c. 1970, Australian) is a leading Australian quantum scientist who is the Director of the University of Queensland Quantum Technology Laboratory.

He is most known for his work in quantum photonics, including the first demonstration of an optical CNOT entangling gate, published in 2003 with Jeremy O'Brien, Geoff Pryde and Tim Ralph, and based on the Knill, Laflamme and Milburn (KLM) protocol and linear optics [170]. He is also very eclectic, having also worked on nuclear physics and marine biology. He's a scientific advisor for Quandela.



Pascale Senellart (1972, French) is a physicist, CNRS research director at the C2N laboratory. She designed and invented a process for manufacturing sources of unique and indistinguishable photons.

These can be used in quantum telecommunications and computing and are made with GaAsAl semiconductor quantum dot trapped in a multi-layered 3D III-V stacked structures, powered by a laser and directly feeding an optical fiber. She co-founded the startup Quandela in 2017 with **Valérian Giesz** (CEO) and **Niccolo Somaschi** (CTO and Chairman) who were a PhD student and a post-doc in her team. Quandela is selling these photon sources and is creating photon qubit-based quantum computers. She is their scientific advisor. Pascale Senellart also launched the Quantum hub of the University Paris-Saclay in November 2019, which brings together public and private research laboratories as well as higher education institutions. She was awarded the CNRS Silver Medal in 2014.



Christine Silberhorn (1974, German) is a researcher and professor working on photon-based quantum computing at the University of Paderborn located between Dortmund and Hanover. She leads there the Integrated Quantum Optics group.

Her laboratory designs and manufactures integrated optronics components, entangled photon sources and quantum array systems. Her team designed a system to convert photon qubits between infrared and visible wavelengths. She also works on optical quantum memories. She was awarded the Leibniz prize in 2011. She cofounded the startup It'sQ in 2022, a quantum photonic computing venture. It turned into a different story, being named QCDESIGN and focusing on qubit control and error correction architecture, and she left the company.



Valentina Parigi (Italian) is an experimental quantum physicist and Professor at Sorbonne Université in Paris, at the Laboratoire Kastler Brossel (LKB).

Her research in the Multimode Quantum Optics group focuses on continuous-variable (CV) quantum optics, multimode quantum states of light, and quantum complex networks. She co-founded Qluster, a startup developing CV photonic quantum computing systems. She got her PhD at the European Laboratory for Non-Linear Spectroscopy (LENS) in Florence and completed post-doctoral work at the Institut d'Optique (LCFIO) and LKB.



Fabio Sciarrino (1978, French Italian) is the director of the Quantum Information Lab at the Sapienza University of Rome and specialized in photonics.

His team is at the origin of many advances in the field, notably in boson sampling experiments. He collaborates with Quandela's team and the C2N of Palaiseau (Pascale Senellart).



Jacqueline Romero (c. 1985, Philippine) is a quantum optics physicist doing research in Australia at the University of Queensland, after completing her PhD in Glasgow, UK.

She is working on optical neuromorphic architectures and on dense encoding of information in photons using several of their characteristics in addition to the usual polarization.



Jelena Vucokic (c. 1975, Serbian) is a research professor at Stanford, working in quantum photonics.

She directs the Nanoscale and Quantum Photonics Lab and the Q-FARM (Quantum Fundamentals, ARchitecture and Machines initiative), an interdisciplinary quantum laboratory. She contributes to developments in

photonics for the development of optical quantum computers. She did her PhD at Caltech in 2002.



Stefanie Barz (c. 1980, German) is a quantum optics professor and researcher at the University of Stuttgart. Her interests include quantum cryptography and quantum telecommunications.

She worked in particular on blind computing with Elham Kashefi and Anne Broadbent. She leads the SiSiQ project funded by the German Ministry of Research with €3.6M of European funding, which aims to create quantum communication infrastructure with silicon photonics.



Sophia Economou (c. 1980, Greek-American) is an Associate Professor in the Department of Physics at Virginia Tech College of Science. She previously worked at the US Naval Research Laboratory.

She is a physicist specialized in the control of quantum dot semiconductor spins and their spin-photon interfaces. She is also a creator of advanced molecular simulation algorithms on quantum computers.

2.4.8 Quantum communications and cryptography



Nicolas Gisin (1952, Swiss) is a physicist specialized in quantum communication. He demonstrated quantum nonlocality in 1997 over a 10 km distance, extending the performance achieved in the laboratory by Alain Aspect in 1982.

He cofounded IDQ in 2001, a startup initially specialized in quantum random number generators using photons passing through a dichroic mirror and then on prepare-and-measure based quantum cryptography and single-photon detectors. It was acquired by SK Telecom in 2018 and then by IonQ in 2025.



Artur Ekert (1961, Polish and English) is a quantum physicist known to be one of the founders of quantum cryptography. He had met Alain Aspect in 1992 to talk to him about this inspiration after discovering the latter's experiments.

This is a fine example of step-by-step inventions, one researcher inspiring another! He was the director of the Singapore Center for Quantum Technology from 2007 to 2020. He is also a teacher at Oxford University and a member of Atos's Scientific Council.



Stephanie Wehner (1977, German) is a physicist working on quantum communication protocols, based at the University of Delft in the Netherlands.

She coordinates the “Quantum Internet Alliance”, one of the projects of the European Quantum Flagship, which plans to deploy a quantum key distribution (QKD) Internet network running in mesh mode. She started her professional life in cybersecurity, detecting system flaws. She is also producing many quantum tech MOOCs.



Eleni Diamanti (1977, Franco-Greek) is a leading specialist and experimenter in the development of photonic resources for quantum cryptography, also working on quantum communication complexity.

She's a CNRS research Director and faculty at LIP6 laboratory from Paris-Sorbonne University. She is the vice-director of the Paris Centre for Quantum Technologies (PCQT) since April 2020. She is also involved in many European projects around quantum key distribution, like the Quantum Internet Alliance and OpenQKD. She is a recipient of a European Research Council Starting Grant. At last, she's a cofounder and a scientific advisor with Julien Laurat for the startup Welinq, created in 2022 with Tom Darras as CEO, which creates cold atom based quantum memories for quantum computer interconnects and quantum repeaters.



Sara Ducci (1971, French) is another teacher-researcher at the same Laboratoire Matériaux et Phénomènes Quantiques (MPQ) where she co-founded in 2002 a team in charge of nonlinear optical devices.

She is working on producing pairs of entangled photons sources based on III-V semiconductors. She is also interested in the characterization (state measurement...) and manipulation of photons. At last, she teaches quantum physics at Ecole Polytechnique.



Sébastien Tanzilli (France) is the director of the InPhyNi physics laboratory in Nice and also the CNRS national quantum program director. He works on quantum cryptography with continuous or discrete keys (CV-QKD and DV-QKD).

He is also specialized in fundamental quantum optics as well as in hybrid quantum systems for the study and realization of quantum communication networks. He was the president of the GDR-IQFA, a community of quantum physics researchers in France (IQFA = Information Quantique, Fondements & Applications) from its creation in 2011 until 2021. It then became GDR TEQ, presided by Anaïs Dréau.



Virginia D'Auria (Italian) is a researcher working on quantum optics transmission systems using continuous and discrete variables and DV/QV hybridization. Having worked at the ENS LKB in Paris, she also worked on photon detectors.

Since 2010, she is part of the photonics group of In-PhyNi and works on discrete and continuous variable quantum communications compatible with optical fibers of telecom operators.



Frédéric Grosshans (1976, French) is a CNRS researcher at LIP6 from Université Paris-Sorbonne, specialized in QKD, repeaters and quantum networks.

He was the creator with Philippe Grangier of the continuous variable QKD. He is also the co-director with Nicolas Treps (from LKB) of the Quantum Information Center Sorbonne of the Alliance Paris-Sorbonne launched in September 2020, which federates quantum research and training of several Parisian quantum groups.

Some other key names here are **Samuel L. Braunstein** and **Stefano Pirandola** (University of York, UK), **Renato Renner** (ETH Zurich), **Dirk Englund** (MIT) and **Julien Laurat** (LKB, WelinQ).

2.4.9 Other academic fields



Jason Alicea (American) is a Professor of Theoretical Physics at Caltech's IQIM (Institute for Quantum Information and Matter). He is a condensed matter physics theorist.

He explores exotic phenomena bridging topology and quantum mechanics like spin liquids and topological phase of matter, which could lead on creating non-Abelian anyons and Majorana fermions qubits. He participates to other fields like cold-atoms physics and quantum measurement.



Alexia Auffèves (1976, French) is a CNRS research director and the director of Singapore's CNRS MajuLab international laboratory since January 2023 after having conducted her research for over 15 years in Grenoble at CNRS Institut Néel.

She started as an experimentalist, doing her PhD thesis creating "Schrodinger's cats" with cold atoms at ENS LKB in Paris with Serge Haroche. She then became a theoretician focused on quantum energetics, with a strong interest in spin-photons interactions and the energetic cost of quantum gates and measurement. She also cocreated the CSM (context, systems, modalities) ontology with Philippe Grangier and the philosopher Nayla Farouki [171]. In the quantum ecosystem, she launched and coordinated QuEnG (Quantum Engineering Grenoble), the Grenoble quantum ecosystem, which became the QuantAlps federation in January 2022. She cofounded the Quantum Energy Initiative in August 2022 with Robert Whitney (a physicist from CNRS LPMMC in Grenoble), Janine Splettstoesser (Chalmers University, Sweden) and Olivier Ezratty. She also runs

with Robert Whitney the Grenoble/Singapore Quantum Energy Team (*|QET*)).



Jacqueline Bloch (1967, French) is a research director at CNRS (PI) in the Centre de Nanosciences et de Nanotechnologies (C2N) lab from CNRS and Université Paris-Saclay, working on polaritons, quasiparticles coupling light and semiconductor matter.

These have potential applications in the creation of quantum simulators based on polariton arrays as well as for quantum metrology. Jacqueline Bloch was awarded the gold medal from CNRS in September 2026.



Alexei Grinbaum (1978, Franco-Russian) is a researcher at CEA-Saclay in Etienne Klein's LARSIM laboratory. He works on the quantum foundations and quantum physics philosophy [172].

He is notably the author of the book "*Les robots et le mal*" (Robots and evil) published in 2018. He is particularly interested in the ethics of science, its acceptance by society and responsible innovation.



Marcus Huber (Austrian) is a research group leader at the IQOQI in Vienna, working on quantum entanglement, qubit state measurement and quantum thermodynamics in general.

In addition to IQOQI, he also worked at the Universities of Bristol, Geneva and Barcelona. He is an advocate of the open research work publication, and behind the creation of the Quantum-Journal.org.



Christiane P. Koch (German) is a professor of theoretical physics at the Free University of Berlin. After her PhD in Germany, she was a post-doc in France and Israel. She is focused on quantum control, with atoms, condensed matter and molecules.

She covers multiple domains including quantum optimal control to shape electromagnetic pulses that drive quantum systems toward their theoretical performance limits, quantum information processing with the design of high-fidelity entangling gates and exploring quantum error correction methods for superconducting qubits, atoms, and molecules. She also works on controlling molecular rotation, exploiting cold molecular collisions, and achieving coherent control of bond-making using lasers and microwave three-wave mixing techniques.



Tommaso Calarco (Italian) is a professor of theoretical physics at Jülich FZ specialized in the field of quantum control. He is also the deus-ex-machina of the European Union quantum scientific strategy since the creation of the EU quantum manifesto in 2016.

2.5 Quantum information science and algorithms creators

Let's continue this long "hall of fame" with some of the main contributors to the creation of quantum information science and algorithms. It is a relatively new discipline that emerged in the early 1980s.

2.5.1 Theory



Paul Benioff (1930-2022, American) proposed in 1979 an early quantum-mechanical model of computation based on a reversible Turing-machine description using 2D lattices of spins $\frac{1}{2}$, based on earlier work from Rolf Landauer on the thermodynamics of computing and Charles Bennett on reversible computing [173].

It was a semi-classical machine concept that didn't yet exploit entanglement and interferences. Paul Benioff was then in a visiting stay at the Centre de Recherche Théorique from CNRS in Marseille, France while being affiliated with the DoE Argonne National Laboratory. He published a second paper in 1982 [174] where he described a theoretical system based on using a two-dimensional lattice of spin $\frac{1}{2}$ systems. Back in the 1980s, the very notion of qubits was not yet in the radar. It appeared much later, in 1995. In Benioff's model, a quantum gate was a Hamiltonian transformation of individual spins that was driven by the Turing quantum machine. David Deutsch extended this in 1985 with the "universal quantum computer" concept.



Yuri Manin (1937-2023, Russian and German) is a mathematician who proposed the idea of creating gate-based quantum computers, in his 1980 book "*Computable and Uncomputable*", then in the USSR.

Then, **Richard Feynman** devised in 1981 the idea of a quantum simulator. Feynman and Benioff were participants of the famous "Physics & Computation" conference in 1981 that was co-organized by IBM and the MIT at the MIT Endicott House [175]. It brought together several well-known scientists in quantum information technology such as Tommaso Toffoli and Edward Fredkin (Figure 59). Rolf Landauer was also among them. It was for this conference that Richard Feynman published his famous paper "Simulating Physics with Computers" which created the concept of quantum simulation [176]. Figure 60 shows a timeline how the concepts of quantum computing came out.



Alexander Holevo (1943, Russian) is a mathematician working in quantum information science and who devised the 1973 Holevo theorem according to which we cannot retrieve more than N bits of useful information from a register of N qubits.

This theorem indirectly validates the fact that it is difficult to do "big data" with a quantum computer in the sense of storing and analyzing large volumes of information. On the other hand, Grover's algorithm makes it possible to quickly find a needle in a haystack, as we will see later. This is the consequence of the wave packet reduction that reduces the qubit state to its basis states $|0\rangle$ and $|1\rangle$ after measurement. He also developed the mathematical basis of quantum communications [178].



David DiVincenzo (1959, American) was a researcher at IBM and the creator of the criteria that define the minimum requirements for a quantum computer with universal gates.

He is now a researcher and professor at RWTH Aachen University in Germany and at the Jülich Research Center. He is a member of the Atos Scientific Council, along with Alain Aspect, Serge Haroche, Artur Ekert and Daniel Esteve, among others.



Gil Kalai (1955, Israeli) is a professor of mathematics at the Hebrew University of Jerusalem and at Yale University. His main ambition is to demonstrate mathematically that it will be impossible to create real universal quantum computers.

It is due to their error rate, even with error correction codes and the notion of logical qubits that assemble physical qubits. He also questioned the reality of the October 2019 Google supremacy performance in several of his writings and conference talks.



Robert Raussendorf (c. 1975, German) is well known for having invented one-way quantum computing and measurement-based quantum computing (MBQC) along with **Hans Briegel** (1962, German) in the early 2000's.

He is an Associate Professor at the Department of Physics and Astronomy of the University of British Columbia. He did his thesis at the Ludwig Maximilians University in Munich, Germany in 2003 on MBQC.



Elham Kashefi (1973, British Iranian) is a research director at CNRS in France, in the LIP6 laboratory from Sorbonne University. She is also the cofounder with Marc Kaplan of VeriQloud, a secure quantum telecommunications startup.

She is a professor of quantum information science at the University of Edinburgh. Originally a mathematician and computer scientist, she became a specialist in quantum communication protocols and quantum algorithms, around topics like code verification and blind quantum computing. She created the BFK blind computing protocol in 2009 with Anne Broadbent and Joe Fitzsimons (who created Horizon Quantum Computing in Singapore).




Figure 59: The participants of the first quantum computing conference in 1981. Source: [177].




Figure 60: A selective conceptual genealogy of quantum computing to remind us that other scientists than Richard Feynman have to be remembered for their contribution. Arrows denote influence. © compilation Olivier Ezratty, 2022.

With her team at LIP6, she is at the origin of the creation of a site on the zoo of quantum communication protocols [179]. And as this was not enough, she is also versed in Quantum Physical Unclonable Functions (QPUF), physical identifiers of quantum and tiltable objects, a topic we briefly cover in this book in the Communications and cryptography part.

In November 2022, Elham Kashefi was appointed as Chief Scientist for NQCC, the UK National Quantum Computing Center, and chairs its Technical Advisory Group.



Anne Broadbent (Canadian) is a mathematician from the University of Ottawa specialized in quantum computing, quantum cryptography and quantum information.

She was a student of Alain Tapp and Gilles Brassard at the Université de Montréal. She created the BFK blind computing protocol in 2009 along with Elham Kashefi and Joe Fitzsimons.



John Watrous (Canadian) is a researcher working at the University of Waterloo, Canada, specialized in quantum algorithms and complexity theory. He demonstrated some complexity classes equivalencies like QIP is in EXP and $\text{QIP}=\text{PSPACE}$.

He also worked on cellular automata. He had previously collaborated with Scott Aaronson. He is the author of the voluminous book “The Theory of Quantum Information” in 2018 [180].

2.5.2 Algorithms



Tommaso Toffoli (1943, Italian then American) is an engineer known for the creation, at the beginning of the 1980s, of the quantum gate bearing his name, a conditional gate with three inputs that is widely used in quantum programming.

After working at MIT, he became a Boston University professor, where he has served since 1995. Like Stephen Wolfram, his interests include cellular automata and artificial life.



Edward Fredkin (1934-2023, American) was a professor at MIT, Boston University and then Carnegie Mellon University. He is the author of the controlled-SWAP gate (CSWAP) that bears his name, which exchanges two target qubits conditioned on the state of a third control qubit. He created the concept of reversible classical computer with Tommaso Toffoli at MIT.

He is also a prolific inventor far beyond quantum computing and is the originator of vehicle identification transponders and automotive geonavigation.

He promotes the notion of “digital philosophy” which reduces the world and its functioning to a giant quantum program, a theory he shares with Seth Lloyd, an idea that has been revived by Elon Musk who believes that the Universe is a gigantic program and that we live in a simulation. Is the “automatic” respect of elementary physical laws a “program”? A thorny philosophical and semantic question!



David Deutsch (1953, Israeli and English) is a physicist from the Quantum Computing Laboratory at Oxford University in the UK.

He devised in 1985 the idea of creating a universal quantum computer using a quantum Turing machine which led him to create in 1989 the gate-based circuits programming model, completing Yuri Manin’s and Paul Benioff’s 1980 ideas [181]. He then described a lot of concepts, including the unitaries used in single qubit gates, and the notion of quantum computing complexity.

In 1989, he published another paper where networks correspond to series of gate operations [182]. Back then, the very name of qubit didn’t exist yet, and was created only in 1995. He is also the author of the first algorithm showing a quantum advantage on an oracle problem, deciding whether a Boolean function is constant or balanced, with two variants, the 1985 single-bit version and the Deutsch-Jozsa generalization to n bits in 1992, co-created with Richard Jozsa. These are promise problems and not search algorithms. At last, he is the originator of “constructor theory”, which he proposed in 2012 and develops with Chiara Marletto.



Umesh Vazirani (Indian American) is a professor at UC Berkeley. He is one of the founders of quantum computing, with his paper co-authored in 1993 with his student Ethan Bernstein, [Quantum Complexity Theory](#).

The Bernstein-Vazirani paper introduced the Fourier sampling technique over $\mathbb{Z}_2^n$ that seeded the quantum Fourier transform (QFT), whose efficient circuit for an arbitrary modulus is due to Don Coppersmith in 1994, the same year it was used by Peter Shor to create his famous integer factoring algorithm that served as a spur to funding research in quantum computing in the USA [183]. The QFT is a founding algorithm used in many other quantum algorithms.



Edward Farhi (1952, American) is a theoretical physicist who has worked in many fields, including high-energy particle physics, particularly at the CERN LHC in Geneva and then at MIT.

He worked with Leonard Susskind on unified theories with electro-weak dynamical symmetry breaking. He and Larry Abbott proposed a model in which quarks, leptons, and massive gauge bosons are composite. He is the creator of adiabatic quantum algorithms and quantum walks. He also introduced with Peter Shor the concept of quantum money in 2010.



Peter Shor (1959, American) is a mathematician who became the father of the algorithm of the same name in 1994 which allows the factorization of integers into prime numbers, based on quantum Fourier transforms (QFT).

He simultaneously created the first quantum discrete-log algorithm (dlog) and, later, the famous nine-qubit flip error and phase error correction algorithm for quantum computers called the “Shor code” [184]. In 1996, he even defined the very principles of fault-tolerant quantum computing [185].

We indirectly owe to him the whole movement of post-quantum cryptography (PQC), generated by the fears created by his factoring algorithm.

PQC is about creating cryptography codes resisting public keys breaking using the Shor algorithm and other quantum algorithms... with quantum computers that do not yet exist. Peter Shor created his famous factorization algorithm while working at Bell Labs. He has been teaching applied mathematics at MIT since 2003.



Daniel R. Simon (American) is the creator of another oracle-based algorithm in 1994, bearing his name. Precisely, his quantum algorithm solves a particular Abelian hidden-subgroup/hidden-period problem (HSP) with an oracle, providing an exponential acceleration compared to classical computing [186].

Daniel Simon was a post-doctoral researcher at the Université de Montréal when he wrote his 1994 paper, and joined Microsoft Research shortly afterwards. He later worked on cybersecurity research until his retirement, always with Microsoft Research.



Lov Grover (1961, Indian-American) is a computer scientist who created the seminal quantum algorithm in 1996 that is said to be a search algorithm in a database but has many more use cases as we will see in the quantum algorithms part of this book.

He got a bachelor's degree from IIT Delhi in 1981 and a PhD from Stanford University in 1984, both in electrical engineering. He worked at the Bell Labs starting from 1984, was a visiting professor at Cornell University, and became an independent researcher in 2008.



Scott Aaronson (1981, American) teaches computer science at the University of Texas at Austin. He is a leading expert in quantum algorithms and complexity theories.

He is notably at the origin of a quantum algorithm used for boson sampling, a way to demonstrate some quantum advantage for photonic based experiments. Bosons are integer spin particles such as photons, while particles such as electrons, neutrons and protons are fermions, with a spin $1/2$.



Dorit Aharonov (1970, Israeli) is a quantum algorithms researcher. She received her PhD in Computer Science in 1999 at the Hebrew University of Jerusalem on "Noisy Quantum Computation" and then did a post-doc at Princeton and Berkeley.

She is credited with the "quantum threshold theorem" co-demonstrated with Michael Ben-Or which states that below a certain error rate threshold, error correction codes can be recursively applied to obtain an arbitrarily low error rate of logical qubits. Threshold theorems in general show scalable suppression of logical error below a certain level of errors under specified noise, locality, control, and architecture assumptions. Correlated or

non-Markovian noise (due to some memory effects) outside those assumptions requires separate analysis.

Dorit Aharonov's uncle is **Yakir Aharonov** (1932, Israeli), a physicist who had worked with David Bohm, among others.



Aram Harrow (American) is a prolific specialist in quantum algorithms. He teaches both quantum physics and quantum computing at MIT. At MIT, he works alongside Peter Shor, Charles Bennett having spent his career at IBM Research in Yorktown Heights.

He is the co-author of the HHL quantum algorithm used to solve linear equations which he created jointly with Avinatan Hasidim and Seth Lloyd [187]. He is also interested in the creation of hybrid classical/quantum algorithms.



Seth Lloyd (1960, American) is a professor at MIT who is a prolific contributor to quantum information and quantum algorithms.

He is the initiator of Quantum Machine Learning, of the concept of qRAM (quantum random access memory), of continuous variables gates-based quantum computing (1999), of quantum radars (2008). He's also the L in the famous HHL quantum linear equation solving algorithm and worked on quantum error correction codes and quantum biology.

In his 2006 book, *Programming the Universe*, Lloyd contends that the universe itself is one big quantum computer producing what we see around us, and ourselves, as it runs a cosmic program. According to Lloyd, once we understand the laws of physics completely, we will be able to use small-scale quantum computing to understand the universe completely as well. In about 600 years.

Seth Lloyd was placed on paid administrative leave by MIT in 2019 and then sanctioned in 2020 with a five-year set of disciplinary measures, because he had not disclosed to his management some Jeffrey Epstein originated funding.

In 2016, he created Turing (2016, USA) with Michele Reilly, a software company working on hybrid classical-NISQ software solutions using AI and quantum machine learning techniques.



Alán Aspuru-Guzik (circa-1978, Mexican-American) is a research director at the University of Toronto, formerly at Harvard, who, among other things, created various quantum chemistry algorithms, a topic we will cover in the section dedicated on quantum algorithms.

He is also a cofounder of Zapata Computing, a startup that developed quantum computing software frameworks, particularly in chemical simulation, and which

ceased operations in October 2024 after a pivot to generative AI and a SPAC listing, before a partial restart in 2026.



Frédéric Magniez (French) is the Director of the CNRS IRIF laboratory mentioned above. He also did run a Chair at Collège de France in Spring 2021.

His research focuses on the design and analysis of randomized algorithms for processing large datasets, as well as the development of quantum computing, particularly algorithms, cryptography and its interactions with physics.

In 2006, he founded and led the national working group for quantum computing, bringing together 20 research groups.



Iordanis Kerenidis (c. 1980, Greek) is a director of research from CNRS at IRIF (Institut de Recherche en Informatique Fondamentale), Paris, France, working on cryptography, quantum communication, quantum complexity theories and quantum machine learning, his latest specialty.

He did his thesis at UC Berkeley under the supervision of Umesh Vazirani, then a post-doc at MIT with Peter Shor where he worked in the same office as Scott Aaronson. He is part of the founding team of QC Ware.

There he leads the R&D in quantum algorithms. He also co-leads the Paris Quantum Ecosystem (PCQC) with Eleni Diamanti. He was one of the members of the parliamentary mission on quantum technologies led by MP Paula Forteza between April 2019 and January 2020.



Ryan Babbush (c. 1989, American) is a Google researcher working on quantum simulation algorithms. His goal is to create commercial quantum chemistry solutions.

In a February 2020 [presentation](#), he did show that chemical simulation with Google's Sycamore 53 qubits processor could not use more than 12 qubits because of its high error rate.



Maria Schuld (c. 1989, German) is a senior researcher and software developer at Xanadu since 2017, based in South Africa at the University of KwaZulu-Natal in Durban, where she got her PhD in quantum machine learning.

She was then a post-doc after a short internship at Microsoft Research in the USA. She is a key contributor to the development of quantum machine learning algorithms, particularly in the field of pattern recognition.



Cristian Calude (1952, Romanian/New Zealander) and **Elena Calude** (Romanian/New Zealander) are researchers in New Zealand, Cristian at the School of Computer Science of the University of Auckland and Elena at the Massey University Albany campus.

They work on quantum algorithms, hybrid quantum algorithms and complexity theories.



Ewin Tang (2000, American) published in July 2018 a paper demonstrating a classical recommendation algorithm nearly as efficient as the gate-based quantum recommendation systems algorithm published by Iordanis Kerenidis and Anupam Prakash in 2016 [\[188\]](#).

They responded by finding a flaw in the reasoning. On close inspection, the quantum algorithm would scale better in some extreme conditions. She was 18 years old at the time. Ewin Tang did her PhD at the University of Washington, was a Miller postdoctoral fellow at UC Berkeley and joined Princeton as an assistant professor in 2026.

We can also mention key researchers like **Andrew Childs** (University of Maryland) **Barbara Terhal** (QuTech, TU Delft), **Garnet Chan** (Caltech), **Ryan Babbush** (Google AI), **Bob Coecke** (Quantinuum), **Zoë Holmes** (EPFL), **Robert Wille** (TUM), **Jens Eisert** (Freie Universität Berlin) and **Frank Wilhelm-Mauch** (Jülich).

2.5.3 Error correction



Michael Freedman (1951, American) is a mathematician who founded and did run the Microsoft Station Q laboratory in Santa Barbara, California. He now works at Google.

He is one of the fathers of topological quantum computing along with Alexei Kitaev. Freedman received the 1986 Fields Medal for his foundational work on topological four-manifolds, including the topological four-dimensional Poincaré conjecture. Grigori Perelman later proved the three-dimensional Poincaré conjecture in 2002-2003.



Alexei Kitaev (1963, Russian and American) is with Michael Freedman one of the fathers of the topological quantum computer concept in 1997, investigated by Microsoft.

He was a researcher at Microsoft Research in the early 2000s and is now working at Caltech and with Google. He has also done a lot of work on error correction codes, including the creation of toric codes, surface codes and

magic states distillation (with **Sergey Bravyi**) and the quantum phase estimation (QPE) algorithm, used in Shor's integer factorization algorithm.



Andrew Steane (1965, English) is a Professor of Physics at Oxford University. He created the so-called Steane quantum error correction code in 1996.

The Steane code encodes one logical qubit in seven physical qubits and corrects any single-qubit error which mix flip and phase errors. Looking at how it works provides good insights on the inner workings of quantum error correction codes, although this particular code will probably not be used when we will have scalable quantum computers.




Daniel Gottesman (1970, American) is a physicist who spent most of his career at the Perimeter Institute in Waterloo, Canada, and moved in 2021 to the University of Maryland.

He did his PhD thesis at Caltech under the supervision of John Preskill. He is known for his work on quantum error correction codes (QEC) and is co-author of the famous Gottesman-Knill's theorem according to which a quantum algorithm using only Clifford gates can be efficiently simulated (meaning, polynomially) on a classical computer. Clifford group quantum gates are based on half and quarter-turn rotations (of the qubit in the Bloch sphere), Hadamard gate and the C-NOT conditional gate. This theorem thus indirectly proves that a basic gate set is insufficient to generate an exponential quantum advantage. Adding a non-Clifford gate such as a T gate or the Toffoli gate to the Clifford set generates an universal gate set that can construct circuits approximating any unitary transformation at a given precision. Non-Clifford resources are necessary for generic quantum advantage.



Mazyar Mirrahimi (circa 1980, Iranian) is a mathematician who moved to quantum physics. He is currently the director of Inria's Quantic laboratory, which specializes in error correction codes and quantum algorithms, among other topics.

He did his post-doc with Michel Devoret at Yale University. Back in 2013, he published a seminal paper on cat-qubits. These are physical qubits using a cavity and a superconducting qubit which exhibit a strong biased noise, with way more phase noise than flip noise. They may reduce logical-qubit overhead under specific hardware and noise assumptions. These cat-qubits are used by the startup Alice & Bob as well as by Amazon, as announced in December 2020.



Zaki Leghtas (Morocco/France) is a researcher based in France in Mazyar Mirrahimi's team and is also specialized in error correction codes and systems. He is notably one of the creators of cat-qubits mentioned above.

These are supposed to enable the creation of logical qubits with fewer than 100 physical qubits. He worked in Michel Devoret's laboratory at Yale University before joining Inria's Quantic team in 2015. He is also affiliated with ENS-PSL and Mines Paris-PSL.

Other names worth mentioning here are **Craig Gidney** (Google AI), **Shruti Puri** (Yale), **Sergey Bravyi** (IBM), **Dave Bacon** (Google), **Liang Jiang** (University of Chicago), **Robert Calderbank** (Duke University), **Gilles Zemor** (University of Bordeaux), **Matthew B. Hastings** and **Nicolas Delfosse** (IonQ), **Jean-Pierre Tillich**, **Anthony Leverrier** and **Christophe Vuillot** (Inria), and **Valentin Savin** (CEA-Leti).

2.5.4 Other domains



Kristel Michielsen (circa-1969, Belgian) is a physicist working at RWTH Aachen University in Germany and at the Jülich Supercomputing Centre (JSC) where she leads the Quantum Information Processing (QIP) research group.

She has contributed to numerous works in quantum computing both in physics and algorithms. She created the QTRL scale, for Quantum Technology Readiness Level, that is used to evaluate the level of maturity of quantum technologies and which we will discuss in the section dedicated to practices in research.



Matthias Troyer (1968, Austrian) is Professor of Computational Physics at ETH Zurich. He joined Microsoft Research in Redmond at the beginning of 2017.

He is one of the creators of the Q# language for quantum programming and of the open-source framework ProjectQ launched in 2016 by ETH Zurich. He is particularly interested in chemical simulation with quantum computers. He received his PhD from ETH Zurich in 1994.



Krysta Svore (c.1978, American) is Vice President, Applied Research in Quantum Computing at Nvidia since November 2025. Beforehand, she was VP of quantum software at Microsoft.

She has a Ph.D. in Computer Science from Columbia University. Her contribution in quantum information science covers a broad range of topics: MBQC, quantum machine learning, contributing to the creation of the LIQUi|> quantum programming language, surface codes, fault-tolerance quantum computing.



Benoît Valiron (1980, France) is a researcher at the CNRS LMF laboratory (Laboratoire Méthodes Formelles) from Université Paris-Saclay and teaching quantum programming and algorithms, including at CentraleSupélec.

This quantum programming specialist is the co-author of the open-source quantum programming language Quipper, which he contributed to create while being at the University of Pennsylvania.



Bettina Heim (c. 1980) was a Microsoft developer specializing in quantum software and responsible for the development of the quantum programming language Q# compiler, promoted by Microsoft since 2017.

It is part of their Quantum Development Kit, currently running on quantum emulators on traditional processors


and now supported on third party hardware proposed on the cloud, including IonQ and Honeywell trapped ion based quantum processors. Since 2023, she is a System Software Engineering Manager at Nvidia.



Cyril Allouche (1975, French) has been leading Atos R&D efforts in Quantum Computing since its beginning in 2015. His team implemented the quantum vision of Thierry Breton, CEO of Atos until 2019.

His work encompasses developing the aQASM (Atos Quantum Assembly Language) quantum programming language and the myQLM quantum programming emulator running on regular personal computers and servers.

Here we are. We've covered a whole lot of people and certainly missed many who should be on this hall of fame list! I'll update it whenever required. We will encounter many of these scientists in this book.



2.5.5 Bibliography and notes

These are the bibliographical references and notes for the **Quantum physics and information researchers** subsection.

- [24] [Ultimate physical limits to computation](#), by Seth Lloyd, arXiv, February 2000 (22 pages).
- [25] [Experimental verification of Landauer's principle linking information and thermodynamics](#), by Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dilenschneider, and Eric Lutz, Nature, March 2012 (3 pages) (preprint on [arXiv](#)).
- [61] [Michio Kaku thinks we'll prove the theory of everything by 2100](#), by Kristin Houser, April 2019.
- [62] [A theory of everything?](#), by Michio Kaku.
- [63] [A new kind of science](#), by Stephen Wolfram, 2002.
- [64] [Finally We May Have a Path to the Fundamental Theory of Physics...and It's Beautiful](#), by Stephen Wolfram, 2020.
- [65] [Some Quantum Mechanical Properties of the Wolfram Model](#), by Jonathan Gorard, April 2020 (66 pages).
- [66] [Book Review: 'A New Kind of Science'](#), by Scott Aaronson, arXiv, July 2002 (14 pages).
- [67] [A New Kind of Science](#), by Stephen Wolfram, 2002.
- [68] [The Autodidactic Universe](#), by Stephon Alexander, Dave Wecker, et al., arXiv, September 2021 (79 pages).
- [69] [Quantum Theory at the Crossroads: Reconsidering the 1927 Solvay Conference](#), by Guido Bacciagaluppi and Antony Valentini, arXiv, October 2009 (553 pages).
- [70] [Électrons et photons, rapports et discussion](#), Institut International de Physique Solvay, July 1928 (307 pages).
- [71] [Maxwell's Demon](#), by R. E. Kastner, arXiv, May 2026 (32 pages).
- [72] [A friendly guide to exorcising Maxwell's demon](#), by A. de Oliveira Junior, Rafael Chaves, et al., arXiv, March 2025 (39 pages).
- [73] [Information and thermodynamics: Experimental verification of Landauer's erasure principle](#), by Antoine Bérut, Artyom Petrosyan, and Sergio Ciliberto, arXiv, March 2015 (26 pages).
- [74] [Maxwell's Demon Is Foiled by the Entropy Cost of Measurement, Not Erasure](#), by R. E. Kastner, arXiv, July 2025 (17 pages).
- [75] [Entangling Disciplines: Causality, Entropy and Time-Travel Paradoxes on a Quantum Computer](#), by Maria Violaris, arXiv, June 2025 (11 pages).
- [76] [Lecture Note on Zeeman effect in Na, Cd, and Hg](#), by Masatsugu Sei Suzuki and Itsuko S. Suzuki, 2011.
- [77] [A new algorithm that simulates the intergalactic medium of the Universe in seconds is developed](#), by the Instituto de Astrofísica de Canarias, May 2022.
- [78] [On a Heuristic Viewpoint Concerning the Production and Transformation of Light](#), by Albert Einstein, 1905 (14 pages).
- [79] [Anti-photon](#), by W. E. Lamb Jr., Applied Physics B, September 1994 (77 pages) .
- [80] [The dramatic story behind general relativity's Nobel Prize snub](#), by Robert Friedman, August 2022.
- [81] [Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor](#), by Mike H. Anderson, Jason R. Ensher, Michael R. Matthews, Carl E. Wieman, and Eric A. Cornell, Science, July 1995 (4 pages) .
- [82] [Bose-Einstein Condensation in a Gas of Sodium Atoms](#), by Kendall B. Davis, Marc-Oliver Mewes, Michael R. Andrews, Nicolaas J. van Druten, Dallin S. Durfee, Dan M. Kurn, and Wolfgang Ketterle, Physical Review Letters, November 1995 (4 pages).
- [83] [When Bose wrote to Einstein: the power of diverse thinking – Physics World](#), by Robert P Crease, February 2024.
- [84] [Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?](#), by Einstein Albert, Boris Podolsky, and Nathan Rosen, Physical Review, March 1935 (4 pages).
- [85] [Einstein attacks quantum theory; Scientists and Two Colleagues Find it is not "complete" even though "correct."](#) by Science Service, New York Times, May 1935.
- [86] [Albert Einstein, David Bohm and Louis de Broglie on the hidden variables of quantum mechanics](#), by Michel Paty, 2007 (29 pages).
- [87] [Einstein and the Quantum - The Quest of the Valiant Swabian](#), by A. Douglas Stone, 2013 (349 pages) .

- [88] [Quantum Model of the Atom](#), by Helen Klus, 2017.
- [89] [In her short life, mathematician Emmy Noether changed the face of physics Noether linked two important concepts in physics: conservation laws and symmetries](#), by Emily Conover, 2018.
- [90] [Women in Science: How Emmy Noether rescued relativity](#), by Robert Lea, Medium, February 2019.
- [91] [The Stern-Gerlach Experiment, Translation of: “Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld”](#), by Martin Bauer, arXiv, January 2023 (5 pages).
- [92] [Wave Properties of Matter and Quantum Mechanics I](#), by Anthony Pitucco, March 2021 (48 slides).
- [93] [Single- and double-slit diffraction of neutrons](#), by Anton Zeilinger et al., Review of Modern Physics, 1988 (7 pages)  (preprint on [ResearchGate](#)).
- [94] [Electron spin and its history](#), by Eugene D. Commins, Annual Review of Nuclear and Particle Science, May 2012 (28 pages)  (preprint on [Uniroma](#)).
- [95] [How Electrons Spin](#), by Charles T. Sebens, arXiv, July 2024 (29 pages).
- [96] [Measurement of the fine-structure constant as a test of the Standard Model](#), by Richard H. Parker, Holger Müller, et al., arXiv, December 2018 (36 pages).
- [97] [Determination of the fine-structure constant with an accuracy of 81 parts per trillion](#), by Léo Morel, Zhibin Yao, Pierre Cladé, and Saïda Guellati-Khélifa, Nature, December 2020 (5 pages)  (preprint on [HAL](#)).
- [98] [Fine-structure constant](#), by Wikipedia contributors, Wikipedia, The Free Encyclopedia accessed 30-April-2026.
- [99] [NIST Team Supersizes ‘Quantum Squeezing’ to Measure Ultrasmall Motion](#), by NIST, 2019.
- [100] [A historical perspective on Copenhagen](#), by David C, 2000 (5 pages).
- [101] [The Present Status of Quantum Mechanics](#), by Erwin Schrödinger, October 1935 (26 pages).
- [102] [Is Schrödinger’s Cat Alive?](#), by Mani L. Bhaumik, arXiv, October 2022 (11 pages).
- [103] [A no-go result on observing quantum superpositions](#), by Guang Ping He, arXiv, May 2023 (7 pages).
- [104] [Quantum Mechanics and Radioactive Disintegration](#), by Ronald Gurney and Edward Condon, Physical Review, February 1929 (14 pages) .
- [105] [Olivia Newton-John’s grandfather Max Born was friend of Albert Einstein](#), by Matthew Alice, 1995.
- [106] [Theory of electrons and positrons](#), by Paul A. Dirac, Nobel Prize site, 1933 (6 pages).
- [107] [Relativistic Effects in Chemistry More Common Than You Thought](#), by Pekka Pyykko, Annual Review of Physical Chemistry, January 2012 (24 pages) .
- [108] [Why is mercury liquid? Or, why do relativistic effects not get into chemistry textbooks?](#), by Lars J, 2018 (4 pages).
- [109] [Relativistic quantum chemistry](#), by Trond Saue, 2019 (110 slides).
- [110] [An introduction to Relativistic Quantum Chemistry](#), by Lucas Visscher, (107 slides).
- [111] [Introduction to Relativistic Quantum Chemistry](#), by Kenneth Dyall and Knut Faegri, 2007 (545 pages).
- [112] [Periodic table of elements](#), by Todd Helmenstine, Science Notes, 2014.
- [113] [Electrons relativists as light sources](#), by Marie-Emmanuelle Couprie, 2011.
- [114] [X-ray diffraction: the basics](#), by Alan Goldman, 2018 (31 slides).
- [115] [Relativistic Quantum Mechanics](#), by David J. Miller, 2008 (116 slides).
- [116] [Von Neumann’s 1927 Trilogy on the Foundations of Quantum Mechanics. Annotated Translations](#), by Anthony Duncan, arXiv, May 2025 (132 pages).
- [117] [The Logic of Quantum Mechanics](#), by Garrett Birkhoff and John Von Neumann, Annals of Mathematics, October 1936 (22 pages).
- [118] [On the Problem of Hidden Variables in Quantum Mechanics](#), by John S. Bell, Reviews of Modern Physics, July 1966 (7 pages) .
- [119] [First Draft of a Report on the EDVAC](#), by John von Neumann, Moore School of Electrical Engineering University of Pennsylvania, June 1945 (51 pages).
- [120] [The Transistor, an Emerging Invention: Bell Labs as a Systems Integrator Rather Than a ‘House of Magic’](#), by Florian Metzler, October 2020 (57 pages).
- [121] [Experimental Test of Parity Conservation in Beta Decay](#), by Chien-Shiung Wu, Ernest Ambler, Raymond W. Hayward, Dale D. Hoppes, and Ralph P. Hudson, Physical Review, February 1957 (3 pages).
- [122] [The Angular Correlation of Scattered Annihilation Radiation](#), by C. S. Wu and I. Shakinov, Physical Review, January 1950 .
- [123] [Quantum mechanics and reality](#), by Bryce S. DeWitt, 1970 (6 pages) .
- [124] [The Many-Worlds Interpretation of Quantum Mechanics](#), by Bryce DeWitt et al., 1973 (146 pages).
- [125] [Experimental realization of Wheeler’s delayed-choice GedankenExperiment](#), by Vincent Jacques, E. Wu, Frédéric Grosshans, François Treussart, Philippe Grangier, Alain Aspect, and Jean-François Roch, arXiv, October 2006 (9 pages).
- [126] [Simulating Physics with Computers](#), by Richard Feynman, May 1981 (22 pages).
- [127] [Quantum Mechanical Computers](#), by Richard Feynman, 1985 (10 pages).
- [128] [Efficiency of Feynman’s quantum computer](#), by Ralph Jason Costales, Alex Gunning, and Tony Dorlas, arXiv, January 2025 (9 pages).
- [129] [On the Einstein-Podolsky-Rosen paradox](#), by John S. Bell, 1964 (6 pages).
- [130] [What Bell Did](#), by Tim Maudlin, arXiv, August 2014 (28 pages).
- [131] [Proposed experiment to test local hidden-variable theories](#), by John Clauser et al., 1969 (5 pages).
- [132] [Nonlocal Hidden-Variable Theories and Quantum Mechanics: An Incompatibility Theorem](#), by Anthony Leggett, Foundations of Physics, October 2003 (25 pages) .
- [133] [The Theory of Local Beables](#), by John S. Bell, Epistemological Letters, 1976 (14 pages).
- [134] [The Two Bell’s Theorems of John Bell](#), by Howard M. Wiseman, arXiv, June 2014 (35 pages).
- [135] [Causarum Investigatio and the Two Bell’s Theorems of John Bell](#), by Howard M. Wiseman and Eric G. Cavalcanti, arXiv, March 2015 (24 pages).
- [136] [A Hall lot of effects](#), Nature Materials, July 2025.
- [137] [Experimental Test of Bell’s Inequalities Using Time-Varying Analyzers](#), by Alain Aspect et al., PRL, December 1982 (4 pages).
- [138] [Experimental tests of Bell’s inequalities: A first-hand account by Alain Aspect](#), by William D. Phillips and Jean Dalibard, arXiv, December 2022 (14 pages).
- [139] [Alain Aspect’s experiments on Bell’s theorem: A turning point in the history of the research on the foundations of quantum mechanics](#), by Olival Freire Junior, arXiv, December 2022 (22 pages).
- [140] [Si Einstein avait su](#), by Alain Aspect, Odile Jacob, December 2025 (360 pages).
- [141] [Experimental loophole-free violation of a Bell inequality using entangled electron spins separated by 1.3 km](#), by B. Hensen, Anaïs Dréau, Ronald Hanson, et al., arXiv, August 2015 (8 pages).

- [142] [A strong loophole-free test of local realism](#), by Lynden K. Shalm, Emanuel Knill, Sae Woo Nam, et al., *Physical Review Letters*, December 2015 (10 pages).
- [143] [Loophole-free Bell inequality violation with superconducting circuits](#), by Simon Storz, Alexandre Blais, Andreas Wallraff, et al., *Nature*, May 2023.
- [144] [Experimental Evidence for a Photon Anticorrelation Effect on a Beam Splitter: A New Light on Single-Photon Interferences](#), by Philippe Grangier, Gérard Roger, and Alain Aspect, *Europhysics Letters*, January 1986 (7 pages).
- [145] [Jean Dalibard, the magneto-optical trap, and the ascent of physics with cold atomic gases](#), by William Phillips, *Académie des Sciences*, 2023 (10 pages).
- [146] [Experimental quantum teleportation](#), by Dik Bouwmeester, Jian-Wei Pan, Klaus Mattle, Manfred Eibl, Harald Weinfurter, and Anton Zeilinger, *Nature*, December 1997 (5 pages)  (preprint on [arXiv](#)).
- [147] [Colloquium: Quantum and Classical Discrete Time Crystals](#), by Michael P. Zaletel, Mikhail Lukin, Christopher Monroe, Chetan Nayak, Frank Wilczek, and Norman Y. Yao, *arXiv*, May 2023 (29 pages).
- [148] [On the interpretation of measurement in quantum theory](#), by H. Dieter Zeh, *Foundations of Physics*, March 1970 (8 pages) .
- [149] [A single quantum cannot be cloned](#), by William K. Wootters and Wojciech H. Zurek, *Nature*, October 1982 (2 pages) .
- [150] [Communication by EPR devices](#), by Dennis Dieks, *Physics Letters A*, November 1982 (2 pages)  (preprint on [dspace.library.uu.nl](#)).
- [151] [Autobiographical Notes of a Physicist](#), by David Mermin, *arXiv*, January 2024 (17 pages).
- [152] [‘Shut up and calculate’ does a disservice to quantum mechanics | Aeon Essays](#), by Jim Baggott, December 2021.
- [153] [Quantum Mechanics, Volume 1: Basic Concepts, Tools, and Applications](#), by Claude Cohen-Tannoudji et al., 2019 (944 pages).
- [154] [Quantum Mechanics, Volume 2: Angular Momentum, Spin, and Approximation Methods](#), by Claude Cohen-Tannoudji et al., 2019 (688 pages).
- [155] [Quantum Mechanics, Volume 3: Fermions, Bosons, Photons, Correlations, and Entanglement](#), by Claude Cohen-Tannoudji et al., 2019 (684 pages).
- [156] [Quantum computing: dream or nightmare?](#), by Serge Haroche et al., *Physics Today*, 1996 (2 pages).
- [157] [Bose-Einstein Condensation of Atomic Hydrogen](#), by Dale G. Fried, Thomas C. Killian, Lorenz Willmann, David Landhuis, Stephen C. Moss, Daniel Kleppner, and Thomas J. Greytak, *Physical Review Letters*, November 1998 (4 pages)  (preprint on [arXiv](#)).
- [158] [High-sensitivity diamond magnetometer with nanoscale resolution](#), by J. M. Taylor, P. Cappellaro, L. Childress, L. Jiang, D. Budker, P. R. Hemmer, A. Yacoby, R. Walsworth, and M. D. Lukin, *Nature Physics*, September 2008 (7 pages).
- [159] [Nanoscale magnetic sensing with an individual electronic spin in diamond](#), by J. R. Maze, P. L. Stanwix, J. S. Hodges, S. Hong, J. M. Taylor, P. Cappellaro, L. Jiang, M. V. Gurudev Dutt, E. Togan, A. S. Zibrov, A. Yacoby, R. L. Walsworth, and M. D. Lukin, *Nature*, October 2008 (4 pages) .
- [160] [Quantum Computations with Cold Trapped Ions](#), by Juan Ignacio Cirac and Peter Zoller, *Physical Review Letters*, May 1995 (4 pages) .
- [161] [Demonstration of a Fundamental Quantum Logic Gate](#), by Christopher Monroe, Dawn M. Meekhof, Brian E. King, Wayne M. Itano, and David J. Wineland, *Physical Review Letters*, December 1995 (4 pages).
- [162] [Integrating a fiber cavity into a wheel trap for strong ion-cavity coupling](#), by Markus Teller, Tracy E. Northup, et al., *arXiv*, July 2022 (9 pages).
- [163] [Manipulating the Quantum State of an Electrical Circuit](#), by Denis Vion, A. Aassime, Audrey Cottet, Philippe Joyez, Hugues Pothier, Cristian Urbina, Daniel Esteve, and Michel H. Devoret, *Science*, May 2002 (4 pages)  (preprint on [arXiv](#)).
- [164] [Implementing Qubits with Superconducting Integrated Circuits](#), by Michel Devoret et al., 2004 (41 pages).
- [165] [Demonstration of Two-Qubit Algorithms with a Superconducting Quantum Processor](#), by Leo DiCarlo et al., *Nature*, June 2009 (9 pages)  (preprint on [arXiv](#)).
- [166] [Classical Reversible Computation by Quantum Coherence](#), by Daniel Loss, *arXiv*, July 2026 (11 pages).
- [167] [Realization of a minimal Kitaev chain in coupled quantum dots](#), by Tom Dvir, Leo P. Kouwenhoven, et al., *Nature*, February 2023.
- [168] [Enhanced Majorana stability in a three-site Kitaev chain](#), by Alberto Bordin, Leo P. Kouwenhoven, Grzegorz P. Mazur, et al., *Nature Nanotechnology*, March 2025.
- [169] [Perspective on Majorana bound-states in hybrid superconductor-semiconductor nanowires](#), by Leo Kouwenhoven, *arXiv*, June 2024 (20 pages).
- [170] [Demonstration of an all-optical quantum controlled-NOT gate](#), by Jeremy L. O’Brien, Geoffrey J. Pryde, Andrew G. White, Timothy C. Ralph, and David Branning, *Nature*, November 2003 (4 pages) .
- [171] [Contexts, Systems and Modalities: a new ontology for quantum mechanics](#), by Alexia Auffèves and Philippe Grangier, *arXiv*, January 2015 (9 pages).
- [172] [Narratives of Quantum Theory in the Age of Quantum Technologies](#), by Alexei Grinbaum, *arXiv*, April 2017 (20 pages).
- [173] [The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines](#), by Paul Benioff, *Journal of Statistical Physics*, June 1979 (30 pages) .
- [174] [Quantum Mechanical Hamiltonian Models of Turing Machines](#), by Paul Benioff, *Journal of Statistical Physics*, October 1981 (32 pages) .
- [175] [How a 1981 conference kickstarted today’s quantum computing era](#), by Harry McCracken, May 2021.
- [176] [Simulating Physics with Computers](#), by Richard Feynman, 1982 (22 pages) received in May 1981.
- [177] [Simulating Physics with Computers](#), by Pinchas Birnbaum and Eran Tromer, December 2005 (28 slides).
- [178] [Alexander S. Holevo’s Researches in Quantum Information Theory in 20th Century](#), by Masahito Hayashi, *arXiv*, April 2024 (17 pages).
- [179] [Protocol Library](#), by Quantum Protocol Zoo, 2026.
- [180] [The Theory of Quantum Information](#), by John Watrous, May 2018 (598 pages).
- [181] [Quantum theory, the Church-Turing principle and the universal quantum computer](#), by David Deutsch, 1985 (21 pages).
- [182] [Quantum computational networks](#), by David Deutsch, *Proceedings of the Royal Society*, September 1989 .
- [183] [An approximate Fourier transform useful in quantum factoring](#), by Don Coppersmith, *arXiv*, December 1994 (8 pages).
- [184] [The Early Days of Quantum Computation](#), by Peter W. Shor, *arXiv*, October 2022 (10 pages).
- [185] [Fault-tolerant quantum computation](#), by Peter W. Shor, *arXiv*, March 1997 (11 pages).
- [186] [On the power of quantum computation](#), by Daniel Simon, 1994 (11 pages).

- [187] [Quantum algorithm for linear systems of equations](#), by Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd, Physical Review Letters, October 2009 (4 pages).
- [188] [A quantum-inspired classical algorithm for recommendation systems](#), by Ewin Tang, arXiv, May 2019 (32 pages).

This **Quantum physics and information researchers** subsection contains 130 bibliographical references and notes containing at least 6,775 pages.

2.6 Research for dummies

As I investigated the broad quantum science and technology landscape, I learned more how fundamental and applied research operate. I did not know much about it before this adventure. Working in the ‘digital world’, as a developer, marketer and in the entrepreneurial ecosystem doesn’t necessarily make you look deeply into the inner workings of research. I discovered many aspects that I am detailing here, particularly with regards to practices, lingua-franca, careers and evaluations.

I went up to publishing a review paper in a peer-reviewed journal, helping me discover all the creation, editing, feedback from and to referees, publishing process and the aftermath with getting contacted by predatory journals and conferences.

If you are a researcher, this is very basic stuff that you already know well. For others, it will clarify some of the vague understanding you might have on how research works. If you have a master’s degree and are looking to get a PhD, it may be useful for your future endeavors.

I would like here to draw a distinction between science, engineering and technology, which are frequently merged to variable degrees in quantum sciences and technologies.

Science builds and organizes open knowledge with testable explanations and predictions about how nature works, including human-made “artificial nature”. It involves understanding the world through observation, experimentation, and theoretical explanations and some mathematical formalism. Scientific methods involve formulating hypotheses, conducting experiments to check it, analyzing data, and drawing conclusions. It seeks to discover underlying principles and laws governing phenomena, aiming for objective, reproducible results. The process may be highly iterative and non-deterministic. Science is frequently funded by the public sector in academic groups.

Engineering applies scientific principles to design, develop, and maintain structures, machines, devices, systems, and processes. It is a creative, iterative process that involves problem-solving, designing, testing, and optimizing. It is inherently interdisciplinary, drawing on knowledge from various scientific fields like physics, mathematics, chemistry, materials science and computer science to create innovative solutions. It considers practical constraints, such as cost, safety, efficiency, environmental impact and competitiveness. Quantum engineering is one domain of engineering. Quantum engineering combines optimization with major feasibility uncertainties, particularly in scaling, integration, reliability, and manufacturability.

Technology is the application of scientific knowledge for practical purposes, particularly in industry. It involves the creation and use of tools, machines, materials, and processes to solve problems or achieve specific goals, with a societal or economic goal and utility. It bridges

the gap between scientific discoveries and their practical applications. In the end, innovation is about putting technology into the hands of users, whether individuals or organizations.

2.6.1 Long-term view

The first key point is the long-term approach in quantum technologies. It can also be found in other branches of physics and so-called deep-tech and hard-tech related sciences. Time scales are measured in decades. It starts with intuition, creativity, passion, rigor, and hard work. These ideas are not always broadly adopted right away. There’s always some resistance with the current scientific establishment.

This long-term history can be observed in condensed matter physics. Brian Josephson devised the Josephson junction in 1962. IBM tried to use it unsuccessfully to build superconducting computers. Anthony Leggett made significant discoveries on macroscopic quantum coherence and dissipation in the early 1980s which led to the creation of the first superconducting qubit in 1999 in Japan [189], then to the quantronium in 2002 [163] and to Google and IBM’s superconducting machines between 2016 and 2020. And we’re not done there since this technology’s scalability has not yet been proven.

Alain Aspect’s work leading to his 2022 Nobel prize in physics award started in the mid-1970s and culminated with his 1982 experiment. Back then, it had no immediate industry application. Fortunately, he was well supported by many laboratories, particularly to build the necessary instrumentation. His work led to the creation of many branches of quantum technology. For example, Artur Ekert was inspired by Alain Aspect’s work to advance the field of quantum cryptography in the early 1990s.

All of this cannot be meticulously planned in advance. Research serendipity must prevail, and even up to some form of educated stupidity [190]. Commercialization comes later, through meetings between specialists from different and complementary disciplines. Innovators are either the researchers themselves, or others, engineers and entrepreneurs, who know how to detect research work having some business potential. Hence the importance of bringing them together in ecosystems of innovation.

However, in its current shape, the quantum startup ecosystem is mostly made of researchers turned into entrepreneurs. This generates its share of misunderstandings with public authorities and policy makers. They are tempted to over-evaluate and measure the performance of basic research, if not to fund it, using only criteria from the business world. On the other hand, and this is particularly true for quantum technologies, research work requires peer reviews. This may give the impression that researchers are both judge and jury. To prevent this from driving decision-makers and people

suspicious, research work must honestly be translated in layman's terms. This should encourage researchers to communicate with broader audiences than their peers. It requires leadership. Scientists must be more involved there, particularly in those times where people are becoming skeptical on science and innovation.

2.6.2 Scientific papers

Fundamental research often aspires to openness. Scientists from all over the world publish their work in scientific papers. Science is about building knowledge and transmitting it. All researchers benefit from the work of other researchers. But papers access varies with publication business models and practices, as well as with security, commercial, and intellectual-property constraints.

Preprints. It is commonplace to see research being published first in open access on the **arXiv** website, created in 1991 at Los Alamos, hosted by Cornell University from 2001 and which became an independent non-profit organization on July 1, 2026, with Penelope Lewis as its first CEO [191]. It hosts about 2.4 million articles and receives around 27,500 new submissions per month in 2026, against 3,000 per month in 2001, with a yearly budget of about \$6.7M. It contains article pre-prints before (and after) they enter a peer reviewing process leading to the paper publication in peer-reviewed journals. The volunteers who manage arXiv still check the basic quality of the submitted papers.




Figure 61: The arXiv home page.

To publish a paper on arXiv, you need to obtain the endorsement of at least one researcher who regularly publishes preprints on arXiv. Preprint publishing enables authors to collect comments on their work. Papers can be easily updated to consider feedback and embed various corrections. Their quantity and quality depend on the author's fame, the topic and the number of researchers who master it [192]. It also depends on the way researchers communicate about their papers in social networks. It allows fast turnaround for debates between scientists [193–195], sometimes in a controversial way [196]. Debates can arise on arXiv with papers commenting a paper, and yet another commenting the comments [197, 198]. This process has even been parodied [199].

arXiv papers can also be withdrawn, for example, if some plagiarism was detected [200] or for some other non documented reason [201–204]. There is even a WithdrawXiv dedicated database for retracted papers

on arXiv, unfortunately with a cut-off date in September 2024 [205]. I have removed several arXiv papers from this book's bibliography after discovering their withdrawal.

arXiv papers must sometimes be taken with a grain of salt, and many have no value at all [206]. Other papers lack critical information like these many quantum algorithms papers with no indication of the supporting hardware configuration and number of qubits [207]. arXiv covered fields contain physics including quantum science, computer science, electrical engineering, mathematics, statistics, information science, economics, quantitative finance and quantitative biology. Computer science is now its largest section by volume, which explains why the AI-generated papers flood described later hit it first.

There are other sites for pre-prints like arXiv, with for example **TechRxiv** on engineering and computer science, **engrXiv** on engineering (with some papers related to quantum technologies), **BioRxiv** in biological science, **ChemRxiv** in chemistry and **medRxiv** in health sciences, **verixiv**, **Cryptology ePrint Archive** for cybersecurity, and **OpenReview.net** are sorts of intermediates between pre-print and peer-reviewed journals. **Preprints.org** is a preprint site from MDPI, a peer-reviewed papers publisher (Entropy, Quantum Reports...). There is even a **ChinaXiv** site, operated by the Chinese Academy of Sciences, which hosts preprints from China.

Peer-reviewed publishing. Between about 9 to 18 months after a preprint, the paper may be published in a peer-reviewed journal. If the delay is too short, it may mean the journal is a predatory one, although delays can be simulated. It is usually published mostly as is, includes some revisions suggested by the “referees” of the review committees, or even with a change of title. In these cases, the version published on arXiv is not necessarily the most recent. It is sometimes updated. The benefits are openness and free access. The peer-reviewed version is usually different from the preprint one. It contains more or less revisions based on the referee's feedback and may even have a new title. Sometimes, papers are directly published in peer-reviewed journals while skipping the preprint stage. It happens frequently with major Google papers and some IBM papers.

Quantum technologies peer-reviewed journals include **Nature** and its various thematic variations like **Nature Communications**, **Science**, **Physical Review X** (PRX), **PRX Quantum**, **Physical Review Research** (PRR), **Physical Review Letters** (PRL), **Physical Review A** (PRA) and **Physical Review B** (PRB), **Quantum Journal**, **Quantum Science and Technology**, **International Journal of Quantum Information**, **Quantum Engineering**, **Advanced Quantum Technologies**, **Quantum Information Processing**, **IEEE Journal of Quantum Electronics**, and **IEEE Transactions on Quantum Engineering** (Figure 62).

In peer-review journals, the reviewers are usually initially unknown to the paper authors. Their identity disclosing depends on the journal's review model: single-anonymous, double-anonymous, transparent, or open. Reviewers provide some feedback on the paper and expect a paper update. The authors provide an updated version and comments that are either accepted or rejected by the reviewer. It can lead authors to modify their claims and even their paper title. When everything's finalized, the paper can be published. Nowadays, the initial paper published on arXiv is also updated to reflect these changes. There is also a special double-blind review process where the authors are unknown from the reviewers to avoid any reviewer bias.

One way to rank scientific publications is to use the Journal Impact Factor (JIF). It indicates how often articles published in the journal are cited. It is the number of citations received during a given year by the items the journal published in the two preceding years, divided by the number of these citable items. A higher Impact Factor suggests greater influence within its field but it doesn't guarantee that your own paper will be highly cited. It is just some statistics and, probably, some significant standard deviation.

Fortunately, in this field, there are only a few predatory journals that have a very poor peer-reviewing process and charge researchers. I am currently spammed by the **Scirea** publishing group to publish my 2023 EPJ-A paper, which doesn't make any sense. You don't publish a paper in two distinct peer-reviewed journals! When an abstract of your talk in a scientific conference is published, you're also quickly spammed by second-grade and predatory journals to publish there your potential upcoming paper.

In most scientific fields, including quantum science, there are many publications but not enough skilled reviewers. This job is sometimes done by PhD students. Sometimes, innovative papers are locked by reviewers, particularly when they are cross disciplined which is frequently the case with quantum science and is a problem when publications are over-segmented.

PhD theses are easier to retrieve and are generally published freely. These are usually good sources of bibliographical information. Beyond the main thesis goal that is to advance science in a usually narrow domain, it generally starts with making an inventory of the state of the art, like in review papers.

Review papers present a state of the art of a field. Their bibliography is generally impressive, sometimes as long as the paper itself. They are a good starting point to study a subject, especially if the paper is not too old. I provide links to many such review papers, particularly on specific qubit types. If the author's pedagogy is good, it can be very useful for learning on your own [208]. A bibliography generally enables you to go deeper into the subject by discovering the need-to-know fundamental texts. There are variations of review papers. Argu-

mentative reviews exhibit a critical examination of the literature, integrative reviews which combine multiple views, historical reviews which cast the evolution of a domain over time, and theoretical reviews which can add your own framework to the domain ([Source](#)).

Perspectives contain individual expert opinions on future directions or key challenges in a given domain. They are usually quite similar to review papers.

A colloquium is a semi-pedagogical article that may be created using content from lectures or formal talks, with a specialized but accessible approach. They are usually "invited papers", written by the author at the request of a peer-reviewed scientific journal.

Authors. Several authors are usually mentioned in scientific papers, up to a very large number. In general, beyond three authors, the first is the one who did the bulk of the work. As illustrated in Figure 63, he/she is usually a PhD student or a post-doc who processed the experience and written a large part of the document, but this may depend on the countries, laboratories and thesis supervisors. The last author is the thesis or research laboratory supervisor. In the latter case, the penultimate author is the thesis director who supervised the work. In between are the other contributors, experimenters or simple reviewers although it does not justify authorship. They may work in a single or in several organizations. It can be a mix of public research laboratories associated with one or several industry vendors. Some coauthors appear very frequently in papers, like Mikhail Lukin (Harvard, USA) and Jens Eisert (Freie Universität Berlin, Germany). It usually comes from the size of their team and the number of collaborations they have with other academic teams.

Some papers have a very large number of authors. It is typical from the papers published by Google Quantum AI, with 77 coauthors for their 2019 supremacy paper and more than 200 for their 2024 error correction paper, which means a large share of their whole team. They probably all contributed to the published work but certainly not equally. The number of authors can be very large for high-energy physics experiment papers. I found one example with 573 authors from 75 research organizations in China and for just 9 pages [210]. But the world record seems to be the Higgs boson CERN paper from 2015 with 5,154 authors, probably nearly all the CERN and partners collaborators in the project [211]. There are more authors than the number of words in the main text and bibliography!

Well crafted papers don't forget to mention the respective contribution of all the authors, like in the example shown in Figure 67. It also mentions reviewers (not those from a peer-review publication), research funding source, any potential competing interest, how the research data can be accessed and the availability of any supplemental material, that is now usually placed at the end of papers in their pre-print format. These supplemental materials can contain technical details and



Figure 62: Some key quantum physics peer-review publications like Nature, Science and Physical Review letters. Nowadays, most of these journal are consumed in their online versions with articles in PDF formats. And, particularly in the field of quantum physics and engineering, preprint versions published on arXiv are used as well without the filter of peer reviewing.

SCIENCE ADVANCES | RESEARCH ARTICLE

PhD student who did the bulk of the work

PHYSICS

Training of quantum circuits on a hybrid quantum computer

D. Zhu^{1*}, N. M. Linke¹, M. Benedetti^{2,3}, K. A. Landsman¹, N. H. Nguyen¹, C. H. Alderete^{1†}, A. Perdomo-Ortiz^{2,4}, N. Korda⁵, A. Garfoot⁵, C. Brecque⁵, L. Egan¹, O. Perdomo⁶, **C. Monroe^{1,7}**

Generative modeling is a flavor of machine learning with applications ranging from computer vision to chemical design. It is expected to be one of the techniques most suited to take advantage of the additional resources provided by near-term quantum computers. Here, we implement a data-driven quantum circuit training algorithm on the canonical Bars-and-Stripes dataset using a quantum-classical hybrid machine. The training proceeds by running parameterized circuits on a trapped ion quantum computer and feeding the results to a classical optimizer. We apply two separate strategies, Particle Swarm and Bayesian optimization to this task. We show that the convergence of the quantum circuit to the target distribution depends critically on both the quantum hardware and classical optimization strategy. Our study represents the first successful training of a high-dimensional universal quantum circuit and highlights the promise and challenges associated with hybrid learning schemes.

other contributors, experimentalists, engineers or simple reviewers

thesis director / principal investigator, or laboratory director who may have contributed to the paper writing or reviewing

Figure 63: Typical presentation of scientific paper's co-authorship. Source: [209].

can be very interesting, like for example, describing the experimental setup and its hardware and/or software engineering.

The other extreme case is a paper having only a single author. It means first that it is probably not a PhD student, otherwise his PhD supervisor would be a coauthor, or the author is a PhD, but he lost the support of his/her supervisor for whatever reason, which is a bad omen and very rare. It may also be a Master student who is looking for a PhD position. You can look at whether the author works in a research institution and his CV.

You can assess the author's network and collaborators if he/she mentions and thanks reviewers or contributors [212]. The author may already be an established one like John Preskill, Peter Shor, Seth Lloyd, Scott Aaronson or Andrew M. Childs [213], so no problem.

Other cases with no attached institution, record or network may mean that the author may be working on some fringe theories in a very isolated fashion, particularly if there is no mention of any help or thanks to anybody

for the research. Sometimes, all the authors have the same last name [214]. It may be a family business!

In many countries, such as the USA, it is common practice to mention authors with the initials of their first and middle names initials. It does not make it easy to search them online, especially for Chinese authors. This is particularly the case when there are many contributors. I try to quote authors with their first name when they are easy to find.




Figure 66: You can configure your Acrobat Reader by adding navigation buttons left and right in your toolbar. It makes it easy to come back to the text after looking at a reference in the bibliography.

- [39] M. V. Burdastyh, S. V. Postolova, T. Proslier, S. S. Ustavshikov, A. V. Antonov, V. M. Vinokur, and A. Y. Mironov, *Scientific Reports* **10**, 1471 (2020).
- [40] C. Müller, J. H. Cole, and J. Lisenfeld, *Reports on Progress in Physics* **82**, 124501 (2019).
- [41] W. A. Phillips, *J. Low Temp. Phys.* **7**, 351 (1972).
- [48] *Relativistic Quantum Mechanics*, by David J, 2008 (116 slides).
- [49] *Von Neumann's 1927 Trilogy on the Foundations of Quantum Mechanics. Annotated Translations*, by Anthony Duncan, arXiv, May 2025.
- [50] *The Logic of Quantum Mechanics*, by Garrett Birkhoff et al, 1936 (22 pages).
- [51] *First Draft of a Report on the EDVAC*, by John von Neumann, Moore School of Electrical Engineering University of Pennsylvania, June 1945 (51 pages).

Figure 64: Why (t.h.) these long bibliographies do not contain any title? And this paper is from 2023 [215]! The reason stems from the fact that space was tight when these papers were printed. Nowadays, they are mostly broadcasted online and space is less of a constraint, let alone the reader's available time.

Figure 65: This is how bibliographical references are presented in this document, putting an emphasis on the paper title, then the key authors, with their first name, then journal, date and number of pages to appreciate the “weight” of the work.




Figure 67: Typical credits at the end of a scientific paper. Source: [216]. This is the typical requirement for some peer-reviewed publications like Nature.

Fringe science preprints. viXra (arXiv in reverse order) is an arXiv for the preprints that will never be published in peer-reviewed publication and are too fringe to be accepted on arXiv (vixra.org/quant). An author only visible on viXra and even not on arXiv is really fringe [217, 218].

Paper mills refers to the industrial-scale production of fraudulent scientific papers, often sold to researchers seeking easy publication for career advancement. These operations undermine the integrity of academic publishing and have become a global concern. They come from fringe organizations often based in China, Russia, Iran, India, and some countries in South-East Asia. They also offer fake peer reviews and guaranteed publications

in second-tier journals, sometimes up to using bribery of journal editors. In biology, these papers reuse existing images. This is a growing phenomenon, although dealing with a rather small share of published papers [219] (Figure 68).

When browsing arXiv new papers from quant-ph every day, I bump from time to time on papers which may be produced by paper mills. They are not yet published in peer-reviewed journals. You can also spot papers where the majority of the content is written with LLM-based chatbots, particularly when they contain a lot of bullet lists of similar sizes. This happens with weak and short review papers with relatively short bibliography and poor analysis. These LLM-generated papers can come

Q This book's bibliography organization

In the first seven editions of this book, bibliographical references were located in document footnotes. I used title convention followed by first author/authors, sometimes their research laboratories or companies, publication date and then number of pages or slides, which help identify immediately the volume and depth of the referenced documents. The many footnotes prevented readers from looking at bibliographical references at the end of the document, which is not practical whether you read a paperback or electronic version of the document.

In the 8th edition of the book, I migrated from Word to L^AT_EX and use a somewhat standard scientific bibliographical style with references at the end of the text. They are however positioned at the end of each key section or subsections, and with a continuous numbering from 1 to over 11,600, which is a rare size for a book bibliography and would maybe deserve to show-up in the Guinness Book of Records. This

makes it easier to look at the references for a particular topic, particularly in the case of the various types of qubits described in this book.

I however kept the references formatting from the Word version with article name with active url on its link, authors (not all of them, at least the first and the last, plus some others that I picked if I happen to know them), the journal, month and year, and article size in number of pages or number of slides for a presentation (Figure 65). It looks better than those cryptic usual bibliographies Figure 64. If you read this document with Acrobat Reader, you can add two buttons in your toolbar to navigate back and forth between different locations in the document as shown in Figure 66.

In this 2026 edition, I added a new feature with the  and  badges associated with paid peer-reviewed journal, the second case pointing to a preprint website like arXiv when available.

from paper mills. Another phenomenon in quantum science that may not relate to paper mills or LLMs are these review papers on quantum computing applications in various domains which start with very low level explanations of what is a qubit, but with a huge gap on quantum algorithms to propose an educated bridge with potential use cases.

Frauds and retractions. It doesn't seem that quantum research is prone to significant paper-milling or even to papers being retracted [220]. On the **RetractionWatch** database, you can find only a few retracted papers in quantum physics, mostly coming from China and India (123 items with “quantum” in the title as of November 2024). It includes the two famous Majorana fermions retractions in Nature, both led from Delft, “Quantized Majorana conductance” in March 2021 and “Epitaxy of advanced nanowire quantum devices” in April 2022, the latter co-signed by the Niels Bohr Institute in Copenhagen [221], plus a third one in Science in November 2022 on chiral Majorana modes. The 2020 Science paper from the Copenhagen group was flagged with an expression of concern but was not retracted.

Another notorious case is the retraction of Ranga Dias's paper on ambient room temperature superconductors by Nature in September 2022 [222–224], followed by a second retraction in November 2023 of his 2023 Nature paper on lutetium hydride. It led the researcher to be dismissed in 2024 by the university that employed him [225]. Sometimes, even preprints in arXiv are removed, but are not accounted for in statistics.




Figure 68: Trend of paper mills production and retraction. The main scientific domain where the phenomenon arises is life sciences. The log scale of the chart still shows that we are not yet invaded by such papers. Source: [219].

2.6.3 Writing a paper

As of September 2026, I had published three papers on arXiv on top of the English editions of this book [226–228], but submitted none of them to a peer-reviewed journal. However, I had obtained a lot of feedback before and after publishing these papers and thank those who helped me craft these review papers at their end. Earlier, in May 2023, came my EPJ-A paper which helped me craft this new part on the tools and methods used to prepare a paper [229].

Papers are usually final or interim reports of a research project led by a PhD or postdoc under the supervision of a thesis director. It can be theoretical, experimental, or mixing both aspects. It must advance the given




Figure 69: The Spyder Python editor that I’ve been using to create several charts in this book, including a qubit fidelities scatter plot. I now use it to create many charts with dynamic data coming directly from my Excel databases, thanks to scripts being generated by generative AIs. Added in 2023, updated in September 2026.

science field in a demonstrative way. Most of the time, the advance is incremental and sometimes, it can be potentially disruptive. Still, you will not find here any advice on how to conduct good research since I am not a researcher nor have a PhD. Others do that well [230]. There are, however, many ethical rules to follow in scientific research and paper writing, for example in the way scientific experimental data is processed and published and on how to avoid plagiarism [231].

I investigate the software technical aspects of publishing. So far, I have always written most of my papers and books using Microsoft Word, with a template and the like, even up to developing some VBA productivity enhancement macros with over 5,000 lines of code, like for creating footnotes, index entries, updating it and the likes.

But the standard in research is $\text{\LaTeX}$ although Word submissions are supported by most publications, and the alternative Typst tool is getting some traction [232]. When you’ve been accustomed to Word or similar tools, getting accustomed to $\text{\LaTeX}$ feels like coming back into the early 1980s, before the advent of “WYSIWYG”, for “what you see is what you get”, when graphical editing interfaces came up with Pagemaker, Word and other software of this period. $\text{\LaTeX}$ is still preferred in the research world since it is a programmable and flexible open-source markup language with a rich software

ecosystem. I adopted it for this edition of the book to follow the tradition in scientific literature. There are many pros and cons. The migration was hard and took me one month. I used Pandoc to convert the Word document to $\text{\LaTeX}$ after cleaning it up with footnotes containing only a single reference. I flattened the tables beforehand. I then created many macros to convert figures into templates (for peoples, companies and countries).

Here are the tools I am using to create $\text{\LaTeX}$ papers and the book, on my Windows laptop:

MiKTeX, a $\text{\LaTeX}$ distribution containing its compilers, that you don’t use directly and is used by your $\text{\LaTeX}$ editor. You install it on your computer before installing a front-end $\text{\LaTeX}$ editor of your choice.

TexStudio, a $\text{\LaTeX}$ editor with some WYSIWYG features that relies on MiKTeX (Figure 72). It helps you edit the .tex file, compile it into a PDF version that you can see on the right of your screen. Clicking on the .tex file finds the corresponding location in PDF and the other way around. One initial problem is to get rid of all the warnings and errors generated by your .tex file. Once you get used to the common type of errors, like forgetting to use } to close some {, it gets better. Two caveats vs Word: PDF generation is quite slow. It now takes over 30 mn in $\text{\LaTeX}$ vs 6 minutes in Word. And the PDF files are 4 times larger. I have to find ways to compress it. You can also improve your

productivity with developing some custom macros for the environment. Those working on a paper collaboratively frequently use the online service **Overleaf** which has a free limited version. Most research organizations and quantum startups have subscriptions for the paid version.

JabRef, a bibliography creation tool, that helps refine the bibliography .bib file associated to your .tex file. For example, it will find various meta data of the incomplete entries in your .bib file, like the DOI URL, all the author names and the like. As a result, the bibliography in your L^AT_EX file will look professional. It can even mention the pages where the articles are mentioned [233]. Another popular tool here is **Zotero**. Both are open-source. Since about 2025, this cleaning task of your bibliography files can also be done using the most recent LLM-based tools as well as with Python scripts generated by these generative AI tools.

Spyder, a Python coding environment to create some scatter plots that I couldn't create with Excel (Figure 69). It works locally on your laptop. I used it to plot qubit fidelities against qubit numbers with different scales, imported from a CSV text file exported from Excel and to showcase the different related vendors. I export the plots in SVG and then PDF files to make sure their rendering is vector based in the book PDF. There are many other such tools like Google Colab.

Several aspects are important in any paper:

Title. It must be concise and catchy, and preferably understandable without being a Ninja in the related domain.

Introduction. This is what 90% of your readers will see from your paper. It must contain all the key messages. It describes the domain, a given problem in the domain, a summary of the work and related results, put into context and a broad perspective on its related impact and applications (Figure 75).

There are other trends in paper titles, like the growth in use of “scalable” (when there’s no proof of it in the paper), “toward” (dream on!), and “advantage” (usually, ill defined) (Figure 71). The volumes are however different, the most preeminent stereotyped word here being “toward”.

You can choose humor to present your paper [234], but may need to avoid it in the title.

Illustrations. This is the second part that is mostly looked at by readers. It frequently contains XY plots or more sophisticated plots. Take care of adding readable and understandable labels in the plot. Illustrations also have a legend that makes the chart and plot easy to understand. All variables must be described, unlike the example in Figure 73 where none of the column labels are explained properly.

Introduction — Quantum computers hold the promise of finding useful applications in combinatorial optimization, quantum chemistry, cryptography decryption, and quantum simulations of many-body quantum systems.

Figure 70: Avoid using the now nearly parodic “*hold the promise*” expression when talking about quantum computing [235, 236]. Find something else! Or release the promise!




Figure 71: Evolution of occurrences of the words advantage, scalable and toward in arXiv papers published since 2018. © Compilation Olivier Ezratty, August 2025. The number of 2025 is the cut-off in August 2025 multiplied by 1.5.

Discussion. It is a wrap up of the paper that lays out the conclusion of the presented work. It should contain key quantitative outcomes from the experiments. See one example of how to structure this part in Figure 74. In the case of quantum algorithms, for example, it would be nice to remind the reader in a simple way what is the quantitative gain in speedup vs classical algorithms. Charts showing this data are usually unreadable by non-expert eyes.

Length. A paper length depends on its goals and contents. Some only have 2 to 4 pages, and others can be much longer. Google’s 2019 supremacy paper was 6 pages long in Nature, with a 66-page supplementary information published on arXiv. Long papers are usually review papers of a given field with many contributing authors.

Bibliography. It is important to source all what you say that belongs to existing knowledge and state-of-the-art in what you cover. It must contain key prior art related to your paper. This is a commonplace work done in the early stages of a thesis. It is good to reference key review papers from your domain.

You usually don’t escape quoting the most famous papers like John Preskill’s 2018 NISQ paper if you write about quantum computing, but adding other papers who add some in-depth perspectives on the most famous papers is also useful for the reader.

Data and code. Research results and datasets can be published in various platforms like **Zenodo**, which was developed under the European OpenAIRE program and




Figure 72: The user interface of TexStudio to edit this book. It can synchronize the source code on the left and the position in the PDF view on the right. I use batch commands to run the full compilation of 24 PDF files for this book, for the A4 and Letter version, full book or its 5 parts, the special files for the Amazon edition and the key takeaways edition. It lasts about one hour on a recent laptop. Collaborative projects use in general the online service Overleaf. One challenge for a large document is to get rid of the many warning messages generated by the compiler (bottom left). 2023-updated in September 2026.

Name	$\mathcal{V}$	$\mathcal{E}$	$\omega(G)$
g136x6	60	1716	48*
g136x8	80	3087	64*
g136x10	100	4827	60*
g136x12	120	6970	96*
g150x6	60	1735	54*

g50634x3	60	1711	30*
g50634x4	80	3066	40*
g50634x5	100	4823	50*
g50634x6	120	6966	60*
1dc.64-c	64	1473	10*
1dc.128-c	128	6657	16*

2dc.128-c	128	2955	5*
1tc.64-c	64	1824	20*
1tc.128-c	128	7616	38*
1et.64-c	64	1752	18*
1et.128-c	128	7456	28*

s3m25x5	125	6892	20*
C125.9	125	6963	34*
hamming6-2	64	1824	32*
hamming6-4	64	704	4*
johnson8-4-4	70	1855	14*
johnson16-2-4	120	5460	8*

Table 5: Benchmarks for quantum algorithms (optimal values are marked with '*').

Figure 73: Example of a bad chart captioning which do not explain what are the values in the columns. Source: [237].

is operated by CERN. The deposits can contain research papers, the experiments data sets, research software, detailed reports and any other digital artefacts.

Using this sort of service is becoming common practice, to make sure experiments are reproducible. Other services like **OSF** (Open Science Framework) also promote open research practices. Otherwise, if the only thing to publish is code, it can be done on **GitHub**.

2.6.4 Publishing a paper

We have already covered how to publish a pre-print on arXiv. How about a peer-reviewed journal? Well, you submit your paper on their website. Each paper provides guidelines and a process like PRX Quantum [239]. I went through a similar process with EPJ-A for my first peer-reviewed paper in 2023.

You submit your paper and can propose a couple referee names. The paper can be rejected right away by




Figure 74: Some recommendations on how to structure the discussion section of a paper. [Source](#).

its editorial team if it doesn't fit the quality standards of the publication. Otherwise, it will be approved for reviewing by some of the referees you have proposed or others, depending on the scope of your paper. It is typical to have about three referees.

After a couple of weeks, the publication sends you the referee's written feedback, which is most of the time anonymous. It can be a pure and simple rejection of the paper, which doesn't meet the reviewer's criteria for publication. Some authors may then try to submit the paper in another less prestigious and/or demanding paper. For example, in the APS journals, you may start with PRX or PRX Quantum, whose impact factors are above 11, then downgrade to PRL (Physical Review Letters, about 8), and after that, to PRR, PRA and PRB, which have a lower "impact factor", in the 2 to 4 range. The shift can even be proposed to the authors by the publishers, with or without an initial negative feedback from referees.

Otherwise you will get feedback recommending the publication of the paper as is (not usual) or conditioned by some revisions (more usual). You respond to this feedback in two ways: with a revised version of your document considering the feedback, with a ready to publish version and a version using colored markups, and a formal and polite point by point response to all the feedback. After the reviewers accept the corrected paper, it will go through some editing process by the publication permanent staff.

They may recommend some changes in figures captioning, make sure all figures are referenced in the text and to order the figures against how they are mentioned in the text. Sometimes, it can also relate to currency, numbering and units' consistency.

Then, your paper is ready for publication. You may need to select your publishing "business model". Basically, you can pick an open-source publishing model, but you are the one who pays for it, for example, \$3,450 for a PRX Quantum paper in 2026, or your research organization already has an agreement with the publisher. The other option is with readers paying the access to your paper either per paper or on a subscription basis and then publishing becomes free for you as an author. You can usually publish on arXiv the accepted version that incorporates the referee's feedback, but not the publisher's typeset "version of record". Other models exist like **Quantum** where the publishing fee is set on a voluntary basis, with a default at 600 euros since January 2025, a reduced fee at 100 euros and a full waiver on request.

This is the process. But nowadays, it is contested by researchers. There's a shortage of reviewers and the growth of published papers is faster than the number of researchers and reviewers. In the case of quantum science, like probably many other scientific fields, there are hundreds of niches and not many experts per niche. It is even harder for "horizontal" papers that may cover a large numbers of domains, from physics to software, science to engineering, and quantum to classical computing. This creates a trend with some researchers

How to construct a *Nature* summary paragraph

basic introduction to the field, comprehensible to a scientist in any discipline (1-2 sentences).

more detailed background, comprehensible to scientists in related disciplines (2-3 sentences)

general problem being addressed by this particular study (1 sentence)

summary of main result (with the words “**here we show**” or their equivalent, in 1 sentence)

what the main result reveals in direct comparison to what was thought to be the case previously, or how the main result adds to previous knowledge (2-3 sentences)

results put into more general context (1-2 sentences)

broader perspective, readily comprehensible to a scientist in any discipline, may be included if the editor considers that the accessibility of the paper is significantly enhanced by their inclusion. (2-3 sentences. Under these circumstances, the length of the paragraph can be up to 300 words. The above example is 190 words without the final section, and 250 words with it.

During cell division, mitotic spindles are assembled by microtubule-based motor proteins^{1,2}. The bipolar organization of spindles is essential for proper segregation of chromosomes, and requires plus-end-directed homotetrameric motor proteins of the widely conserved kinesin-5 (BimC) family³. Hypotheses for bipolar spindle formation include the 'push-pull mitotic muscle' model, in which kinesin-5 and opposing motor proteins act between overlapping microtubules^{2,4,5}. However, the precise roles of kinesin-5 during this process are unknown. Here we show that the vertebrate kinesin-5 Eg5 drives the sliding of microtubules depending on their relative orientation. We found in controlled *in vitro* assays that Eg5 has the remarkable capability of simultaneously moving at $\sim 20 \text{ nm s}^{-1}$ towards the plus-ends of each of the two microtubules it crosslinks. For anti-parallel microtubules, this results in relative sliding at $\sim 40 \text{ nm s}^{-1}$, comparable to spindle pole separation rates *in vivo*⁶. Furthermore, we found that Eg5 can tether microtubule plus-ends, suggesting an additional microtubule-binding mode for Eg5. Our results demonstrate how members of the kinesin-5 family are likely to function in mitosis, pushing apart interpolar microtubules as well as recruiting microtubules into bundles that are subsequently polarized by relative sliding. We anticipate our assay to be a starting point for more sophisticated *in vitro* models of mitotic spindles. For example, the individual and combined action of multiple mitotic motors could be tested, including minus-end-directed motors opposing Eg5 motility. Furthermore, Eg5 inhibition is a major target of anti-cancer drug development, and a well-defined and quantitative assay for motor function will be relevant for such developments.

Figure 75: Nature introduction guidelines. Source: [238].

choosing to publish preprints on sites like arXiv and not on peer-review papers [240].

The reviewing process can also be plagued with bad reviewers behavior. I don't know how frequently it happens. There's been cases when they were forcing the authors to cite their work even when it was irrelevant to the paper (see Figure 76 and [241] in the Acknowledgements section).

Researchers can organize their own review process with sending their draft papers to other researchers they know, or do not know, and get their feedback. These appear in the acknowledgment section from their papers. Doing this well can gather more feedback than from paper referees. But even in that situation, friendly researcher's available time is scarce. So you end-up fighting for attention with your peers. At some point, your friendly preprint reviewers can become coauthors if they add some content to the paper. Some researchers appreciate the offer since it may improve their h-index, some not, when they do not agree on the whole text content.




Figure 76: Example of referees asking authors to add references of their own work which is irrelevant with regards to the paper content. In that particular case, it led to the forced retraction of the paper by the publication and launching a new reviewing process with different reviewers [242].

2.6.5 Communicating on a paper

Nowadays, it is a common practice for researchers to communicate on their preprint or peer-reviewed papers

in social networks. I see such news on Twitter (now X) and LinkedIn. When you follow many quantum physics or quantum computing researchers, you get good “head-up” news about interesting papers. Also, some researchers make an effort to summarize the key points of their papers and ask for comments. A recent survey did show that promoting your paper on X improved the chances of the paper being cited in other papers [243].

Scientific publications can be discovered by following the RSS feeds of arXiv, reference specialized papers, in addition, from scientific news feeds of online media or popular scientific press. I also discover new interesting papers with scanning scientific conferences presentations.

In the case of quantum technologies, the “tech” media often broadcasts scientific news dressed-up with sensationalism and exaggerations. This often stems from the propensity of laboratory communicators or sometimes researchers themselves to make shortcuts between their work and its potential usage that may be very long-term. Sometimes, the communication on the paper uses a title that is totally unrelated to the underlying scientific work [244].

It is even stronger when communication comes from a large company such as Google or when the article was written by the laboratory’s communication branch (example in Figure 77). The job of the technology screener is about sorting this out. When your local non-English speaking media broadcasts such information, it is often necessary to start by identifying the original paper which is possibly quoted at the end of the article. Sometimes, you discover blatant translation errors that entirely twist the scope of the covered scientific advance.

2.6.6 Predatory, scam and fake conferences

There are really many conferences and workshops organized throughout the world on the myriads of different special domains of interest in quantum science and technologies. We provide a list of the key ones in the appendix of this book. The conferences and workshops have different tracks; keynotes and tutorials delivered by invited speakers, sometimes panels, and poster sessions and presentations, open to PhD students and postdocs.

When you publish your first papers in peer-reviewed journals, you discover another artefact from this ecosystem: predatory and fake conferences [248]. They invite you as pre-selected as “invited speakers”. These conferences are predatory when they are imitating an existing conference with a slightly different name. But they take place in weird locations, most of the time, second-grade hotels. They charge you as invited speakers when usually you are not charged, and the most prestigious ones will even cover the travel expenses of invited speakers.

These conferences can even be entirely fake, meaning, they collect the money from invited speakers and vanish in a parallel world with your money.

You can spot the weirdness of these events by multiple means:

- Looking at the invited speakers and organizing committee pedigree and countries of origin. Sometimes, the advertised invited speakers have not even been contacted nor accepted any participation.
- The thumbnail pictures of the speakers do not even respect height and width ratios ([example](#)).
- The speakers are too prestigious to be credible like with the [Global Summit on Quantum Computing and Engineering](#) which advertises presentations from John Preskill and Peter Shor, and... Richard Feynman, on top of having IBM, Microsoft, Google, PsiQuantum, Quantinuum as sponsors. You will never see them altogether!
- The invited speakers have no experience or skill in the main topic of the conference. For example, in the [Global Summit on Quantum Computing](#) organized in Barcelona in March 2026, most of the announced plenary and keynote speakers are not related to quantum computing, but rather to quantum materials. It is organized by ScienceQuest from India.
- Such “global conferences” sometimes exhibit past editions pictures which shows a class of about 20 participants ([example](#)).

So, beware!

2.6.7 Papers analysis and classification

I have been updating this book on a yearly basis since 2018. Starting in 2020, I used arXiv and scientific publications are key sources for the updates. As a general rule, when I discover the existence of an article, I search for it on Google Search with the name followed by “file-type:PDF” and I find it free of charge in more than 90% of the cases on arXiv or on the ResearchGate site, the researchers’ reference social network. I also now use the **Unpaywall** Chrome extension which automatically finds an open access PDF version of the paper when you are on a journal web site. It either opens the PDF of the published version when access is open or automatically finds the equivalent arXiv version if not.

The bulk of my work is about classifying the collected information: what is it about and how does it fit into the web of quantum technologies? Until 2025, no artificial intelligence could automate this process. Frontier models can now do a large part of the classification work, but the mapping to the structure of this book still requires a human in the loop. This classification task is a tedious one and you can be easily misled with reading a paper title or press release too quickly. Here is one interesting example with a news post wrongly describing the European LSQuaNT project as an initiative to provide quantum computing solutions to the transportation industry. It is a project related to fundamental quantum

Q Reviewer's instructions

Here are the instructions for reviewers from the Royal Society Publishing for which I had the occasion to review a paper in 2026 [245]:

- *Writing style and appropriateness for a wide audience.*
- *Scientific accuracy, including statistical analysis. Referees may suggest that the article be sent to a specialist statistical reviewer. [...]*
- *Whether the research methods are appropriate, and evidence is provided for the conclusions drawn.*
- *Use of suitable illustrations, tables and supplementary material to illustrate results.*
- *Appropriate length. Each article should be of the shortest length required to contain all useful and relevant information, and no longer.*
- *References should be appropriate, relevant, and devoid of unnecessary self-citations.*
- *Data sharing. It is the policy of the journal that authors make any data, code and research materials supporting the results in the article publicly available on publication, either in a repository or as supplementary material, and appropriately cite this in the reference list [...].*
- *Transparency of information. Papers must include appropriate end sections: authors' contributions, competing interests, ethics statement (where relevant), data accessibility and funding. Please note in your comments if you feel that anything is missing or concerns you.*
- *Please alert the Editorial Office if you suspect inappropriate use of AI, for example unusual wording or incorrect references. [...]*
- *Electronic supplementary material. Supplementary material should be reviewed in addition to the main text. Please note that the main article as published should stand on its own merit.*

MARCH 18, 2021

Scientists take step towards quantum supremacy

by National University of Science and Technology MISIS



Qubit production process. Credit: Sergey Gnuskov/NUST MISIS

A Russian-German research team has created a quantum sensor that grants access to measurement and manipulation of individual two-level defects in qubits. The study by NUST MISIS, Russian Quantum Center and the Karlsruhe Institute of

this would mean they are building some sort of quantum computer...

... but it's just about a new sensor measuring the quality of superconducting qubits using some new materials

Article | Open Access | Published: 05 February 2021

Quantum sensors for microscopic tunneling systems

Alexander Bilmes, Serhii Volosheniuk, Jan David Brehm, Alexey V. Ustinov & Jürgen Lisenfeld

npj Quantum Information 7, Article number: 27 (2021) | Cite this article

836 Accesses | 8 Altmetric | Metrics




Fig. 1 Experimental setup and qubit sample. **a** Schematic of the transmon qubit circuit to study TLS in deposited materials. The qubit island (red) is connected to ground by an additional small capacitor containing the material to be studied. **b** Setup for tuning TLS by applied mechanical strain and a global DC-electric field. **c** Photograph of the qubit island. **d** DC-SQUID connecting the qubit island to ground, and a zoom onto one of the small Josephson junctions. The large-area stray junctions are highlighted in light green. **e** Sample capacitor in overlap geometry as used in this work. It employs a 50-nm thick layer of AlO_x as the sample dielectric. **f** An alternative sample capacitor design consists of two coplanar electrodes separated by a so-called nanogap of a few tens of nm. Here, the sample material can be deposited in a last fabrication step (see inset), or the nanogap can be left uncovered to study individual TLS in native surface oxides.

Figure 77: Example of a scientific paper presented with outrageous claims by its lab communication department. Sources: [246] and [247].

physics and digital simulation of quantum transport, a condensed matter phenomenon [249].

What is the actual progress made with regards to the state of the art? You can rely on these typical recommendations: read the introduction and not just the abstract, identify the problem that the authors are trying to solve and how they are advancing the state of the art, look

at the data and charts, interpret their meaning, and identify any missing data. Also, read the conclusion. If you can't decipher the paper content, do a search of other more generalist web sites mentioning it.

The most efficient way to analyze a paper is to discuss it with a domain specialist when you know one. She/he may not have yet read the paper you have bumped into,

so patience may be required. You may also benefit from being in touch with the large academic research groups which have the habit of analyzing collectively and in detail the major papers published in their domain. How can you check a paper when the specialists in your own network have not yet had the time to do so? You then either need to be patient or to do it on your own. For big news related to quantum computing, one can wait for the next post from Scott Aaronson or a laconic tweet from John Preskill and other luminaries.

In general, a paper presenting a breakthrough that will allow a quantum computer to operate at room temperature or ahead of all others is often in reality presenting simple incremental evolution in the development of a particular type of qubit or a “very low TRL” exotic quantum object that is even not yet a controllable qubit. After reading it, it looks like your tiny hairy dog after the shower! In many cases, quantum science-related papers are inaccessible, requiring solid mathematical and/or physics background. Even quantum science specialists have a hard time interpreting many papers.

You also frequently come across a set of Russian dolls concepts with unknown concepts referring to other unknown concepts, and so on, with numerous bibliographical references being presented as some knowledge given on your side. This is some sort of involuntary humor of scientific complexity. Hopefully, some papers do not use too much jargon and manage to deal with a big fundamental question by making it understandable to many specialists in their discipline and well beyond. You can exercise your reading and analysis skills on very famous papers like Mikhail Lukin’s 2023 paper on logical qubits in neutral atom arrays, Google’s 2019 supremacy paper and their 2024 demonstration of quantum error correction below the surface code threshold. You also find weaknesses randomly in other papers dealing with the subject you are investigating.

Finally, I use arXiv as soon as I come across a startup that defines in too broad terms what it does without any technology specifics. It is so commonplace now! A search starts with finding the startup scientific founder, then with identifying their research work that they are probably willing to package in their freshly created startup. In their bibliography creation process, researchers also look at **Google Scholar** and on **SciRate**, where discussions take place around pre-print papers published on arXiv, and the new **alphaXiv** which does the same since 2024. You can also do some searches on patents.

Here are some interesting examples of papers whose title refers to several cryptic concepts: “The Franke-Gorini-Kossakowski-Lindblad-Sudarshan (FGKLS) Equation for Two-Dimensional Systems” [250], “Floquet integrability and long-range entanglement generation in the one-dimensional quantum Potts model” [251], “Probing Lorentz-Invariance-Violation Induced Nonthermal Unruh Effect in Quasi-Two-Dimensional Dipolar Condensates” [252], “Emergent quantum mechanics of the event-universe, quantization of events via Denrographic

Hologram Theory” [253] and “Emergent Sasaki-Einstein geometry and AdS/CFT” [254] which has some connections with “Exploring uberholography” [255] which deals with some quantum error correction code. To some extent, this complexity can be fun. See also “Variational quantum algorithm for measurement extraction from the Navier-Stokes, Einstein, Maxwell, B-type, Lin-Tsien, Camassa-Holm, DSW, H-S, KdV-B, non-homogeneous KdV, generalized KdV, KdV, translational KdV, sKdV, B-L and Airy equations” [256] which requires a significant mathematical background, “Quantum Pontryagin Neural Networks in Gamkrelidze Form Subjected to the Purity of Quantum Channels” [257] and “Multi-self-loop Lackadaisical Quantum Walk with Partial Phase Inversion” [258].

How to sort and rank scientific papers? You can spot the authors and if the list includes some preeminent principal investigator, it will be a good sign. In the opposite, a single author with no affiliation [259], too many self-references in the bibliography, and no thank you and acknowledgement to third party scientists is usually a bad sign. Successful lone inventors and scientists are rare in this quantum world. You can also look at the presented figures of merit and find conclusions that are frequently not as flashy as in the introduction, particularly with quantum computing algorithms and case studies. You can also spot whether an arXiv preprint has been later published in a peer-reviewed journal using the DOI link on the arXiv page. The peer-review version may be improved. However, a peer-reviewed paper is not necessarily rock-solid. It is becoming more difficult for these journals to find skilled reviewers, particularly for complex multidisciplinary papers.

In the case of quantum computing, I have noticed that the figures of merit have changed over time. As the systems grow in complexity, there are more benchmarks being used that make comparisons difficult. For example, you’ll notice that quantum processors are tested with Bell states or GHZ multi-qubit states instead of just single and two-qubit gates. You then need to figure out the meaning of all these numbers. You may bump into some papers which are totally wrong scientifically, but not criticized by any scientists. It can happen when, for example, someone collapses the complexity classes and asserts that NP-complete (the hardest problems in NP, those to which every polynomially verifiable problem can be reduced in polynomial time) is contained in BQP (problems solvable in polynomial time with bounded error on a quantum computer) [260]. They don’t have the time and want to avoid the Streisand effect, which would amplify the bad buzz.

We must recognize our limits and understand that we are not protected from believing scientific hoaxes like the famous one created by **Alan D. Sokal** in 1996. It merged social sciences and quantum gravity and was published in a social science publication, not a quantum physics one [261]. Also, take care about papers published on April fool’s day, or enjoy them [262, 263]!

Hopefully, quantum scientific publications are way more serious than most of the quantum hype that is conveyed by general news with their amazing amplification capabilities. You will read time and again that quantum computing will drive autonomous cars, create quantum intelligent robots, reduce CO₂ emissions, cure cancers, help Tesla or Hyundai build top-notch batteries or that quantum communications will teleport your data faster than light around the Earth.

Most of these assertions will flourish when the IBMs and Googles of this world make fancy announcements or after your government launches its own “billion dollars” national quantum plan. But they are at least unproven if not entirely false. Who’s going to reveal it to you?

2.6.8 The generative AI revolution

The generative AI frenzy started with ChatGPT and the LLM (large language models) technology buzz since late 2022. These tools progressed at an incredible pace. What was true in the early years (2022-2023) is not valid anymore.

Early LLMs usage

It is no surprise that the quantum ecosystem was curious about it. Classical machine learning was already used to conduct quantum research in many areas, like for the tuning of the frequencies used in tunable frequencies superconducting qubits with Google, in quantum error mitigation, or in designing complicated quantum photonics experiments [264–266].

But LLMs puts us in a different world, when nearly all cognitive and creative processes can get some help from generative AI. A misplaced help was about predicting that quantum computing would make it easier to build large LLMs. It will probably not happen since quantum computing is not relevant for big data applications as we’ll uncover in this book.

When you know how ChatGPT and its equivalents are constructed and their limits, you can still use them daily as an extension to your regular search tools. They can help you segment topics, differentiate two algorithms, list characteristics of a system and so on. Initially, it worked better with intemporal science since LLMs were not frequently updated [267]. But since 2025, they are able to analyze content collected on the web.

Other use cases are created like using ChatGPT to construct new circuits like variational algorithms ansatzes (we cover these notions later in the book) [268] and even design quantum experiments [269]. Indeed, once all the existing formalized knowledge from a scientific field is used to train an LLM, the possibilities are limitless. A GPT LLM can for example learn the measurement outcomes of cold-atom qubits [270]. When used as a giant knowledge repository, there’s not much doubt that these tools can help advance science [271]. You need to understand that these tools reasoning capabilities

are limited although they are continuously improving thanks to the use of “chain of thoughts” and reinforcement learning by human feedback (RLHF) [272–275]. These are not anymore just tools trying to predict the next word in a conversation!

Google **NotebookLM** (now integrated in Gemini) can be very useful to analyze documents, including this book. It can even create very thorough captions for the cryptic figures of many papers. Microsoft CoPilot does this right away with a link on a paper. It could even fact-check some tables I prepared for this book, find the information that was OK, the wrong data, and so on. It makes some mistakes but if you spot them, you still gain a lot of time.

However, when used as good assistants, these LLM-based chatbots are very useful to test conjectures and hypothesis, consolidate state of the art in specific domains, even prove some theorems. They may not have real reasoning capabilities and only superficial ones, but the ability to consolidate and mash-up existing knowledge and reasoning is already providing huge value. In one recent example, a quantum physicist used GPT-5 to craft a paper on quantum field theory [276]. You can also fact-check a lot of writings, data and tables with these tools. Other basic things are amazing like uploading the image version of a table and getting the $\text{\LaTeX}$ version of the table in response, with no errors, even in the mathematical expressions conversion.

Then, for this 2026 edition, I switched to using mostly Claude Pro (for advanced stuff), ChatGPT (for fact-checking), and Gemini Pro (for more regular and less demanding tasks). I describe these in the making of of this book at its very end.

Side effects

There are of course some side effects of the abuse of these chatbots, when some scientists or students overly rely on them not only to rephrase their papers but also to entirely write it. It shows up in these long lists with the same number of words as well as in the empathetic writing style [277]. Some tools detect AI-generated text like Grammarly and GPTZero but they are themselves generating many false positives and false negatives [278].

Some early ChatGPT experiments were sending mixed signals about LLMs. In 2023, Chris Ferrie was picky when he denounced the inexactitudes of ChatGPT trying to explain what quantum computing is [279], when ChatGPT was just parroting what is being repeated ad nauseam in the literature. In 2023 as well, Scott Aaronson, who worked for a while at OpenAI, had GPT-4 take his final course exam, and it fared relatively well with a 69/100 grade when his students got an average of 74.4/100 [280]. I tested OpenAI ChatGPT 3.5 and Google Bard to get answers to the questions I laid out starting in the Why? part of the book (1), and they fared relatively well although not responding with many details or supporting data and links. But ChatGPT 4.0

fared much better and seemed to generate fairly good results.

A 2023 survey conducted by Nature among more than 1,600 researchers did even show that over 25% of them were using LLMs to help write manuscripts and about 15%, one in six, to write responses to research grant RFPs [281]. The most commonly detected abuses are leftover “Regenerate response” or “as of my knowledge cut-off in (date)” in papers [282–284]. This survey included researchers in all disciplines with an apparent LLM usage peak with computer science researchers. We could expect that quantum physics researchers would be less keen than quantum information science researchers to use LLMs in writing RFP responses. ChatGPT can lead to some abuses in scientific paper writing. They are sometimes not detected by peer-reviewing processes [285].

Research work is now abundant to see whether LLM-based chatbots are affecting learners performance in one way or the other [286], generating some cognitive debt [287]. The results are mixed given the measurement tools are not really standardized and surveys are frequently done on a narrow domain that can’t be generalized.

A worrying side effect is the surge of hallucinated references generated by LLMs [288] (Figure 78). Fortunately, this phenomenon can be avoided. Of course, with checking each reference provided by a LLM, and with automating the process. LLMs like Claude Pro and ChatGPT 5.6 and beyond seem to be now immune to this effect, although they can still generate existing reference, but which are remote from the insertion point topic in your writings.

Another identified side effect is the way the “slop” language generated by LLMs influences human written and spoken communication. It shows up in the growth of usage of some terms like “delving” and “crucial” [277, 289]. You can also see the effect at play in the flood of nonsense posts on LinkedIn [290].

Finally, LLM-based harassment is emerging. It happens when a reader runs a newly published paper through an LLM and then either flags the output on social networks or emails it to the authors, usually without having verified any of it. There’s an asymmetry here between the cost of the fact-checker who can obtain some output from an LLM in a couple minutes or so, and the burden for the recipient who may need to check his content for hours. It can also show up with letters sent to peer-reviewed journal editors. Some good practices here include, for the fact-checkers, making the claim specific and verifiable, with an equation, a figure or a page number, to contact the authors privately before going public, to disclose which model was used. On their side, the authors can also run an LLM-based check on the claim to save time. One possible preemption for authors is to fact-check their past publications with LLMs. Many academics who did that found mistakes in

their own papers [291]. Now, will they correct it, either on arXiv or in the peer-review journals when published there?




Figure 78: The growth of hallucinated references in papers published on arXiv. The data is based on auditing 111 million references in 2.5 million papers published in arXiv, bioRxiv, SSRN, and PubMed Central. About 146,932 hallucinated citations were found in 2025 papers. Source: [288].

In a randomized controlled study run over more than 20,000 ICLR 2025 reviews, LLM feedback sent to the reviewers produced more specific reviews and longer author-reviewer discussions, with 27% of the reviewers who received it updating their review [292].

Living with the revolution

I started to explore the topic early on in 2024 [293] and was then wondering how science would operate once nearly all human knowledge becomes easy to exploit. I started to use LLMs to update this book in 2023. My usage grew every year as these tools made so much progress.

2026 was a defining year, showing the extent of how generative AI and LLM-based tools could revolutionize science at all levels: fundamental research, applied research, engineering, teaching, studying, writing, coding, evaluating third-party and your own papers, do fact-checking of anything, and even implement reasoning at amazing levels. We are way beyond about whether AI passes the Turing test. AI is now your teacher, your student, your collaborator, your coach and to some extent your Swiss-knife assistant.

LLMs based tools are now not just predicting the next word in some text. They are now complemented by Reasoning language models (RLMs) [294] and agentic reasoning [295]. LLMs are now able to implement system 1 quick reasoning and even sometimes system 2 thoughtful reasoning, up to extend the field of scientific discovery [296–301] and their limits are of course investigated [302–305]. Even though it may not reach the sophistication of some scientific reasoning, the brute force applied to a large knowledge base appears like being some form of advance reasoning. All this boils down to understanding the role of language in reasoning.

The LLM revolution started to hit the field of mathematics. In April 2026, Terence Tao demonstrated a lemma from one theorem with the help from ChatGPT [306]. In July 2026, he conjectured that “*AI tools will, reasonably soon, become capable of performing a reasonable fraction of research-level mathematical tasks, with reasonable levels of success, quality, supervision, and cost*” [307, 308]. Then, a flood of conjectures were proved or disproved, including 10 done by an internal OpenAI model [309, 310], which was released in September 2026 as “Astra”. One of these proofs got a correction [311]. Others like Max Weinreich are totally opposed to the use of LLMs for creating mathematical proof [312].

But other fields of science are impacted as well, including quantum physics [313], finding quantum algorithms suitable for specific quantum computing hardware [314], and quantum error correction [315]. Some researchers are even toying with the idea to have LLMs reinvent scientific discoveries, including the key ones related to quantum physics, based on the existing knowledge cut-off of their creation [316].

As of 2026, LLMs are good for deductive reasoning [317], analogical reasoning [318] and average deductive reasoning [317]. They are yet not faring well with inductive reasoning [319]. The largest difficulty is with abductive reasoning, connecting the physical world and their mathematical models [320]. But progress is so fast that you should never say never.

As arXiv started to be flooded with low-value papers generated by AI, particularly in the computer science section, the arXiv team sent a chilling warning. On May 14, 2026, Thomas G. Dietterich, Professor Emeritus at Oregon State University and chair of arXiv’s computer science section moderators, published on X a warning to arXiv authors [321]: “*Our Code of Conduct states that by signing your name as an author of a paper, each author takes full responsibility for all its contents, irrespective of how the contents were generated. If generative AI tools generate inappropriate language, plagiarized content, biased content, errors, mistakes, incorrect references, or misleading content, and that output is included in scientific works, it is the responsibility of the author(s). We have recently clarified our penalties for this. If a submission contains incontrovertible evidence that the authors did not check the results of LLM generation, this means we can’t trust anything in the paper. The penalty is a 1-year ban from arXiv followed by the requirement that subsequent arXiv submissions must first be accepted at a reputable peer-reviewed venue. Examples of incontrovertible evidence: hallucinated references, meta-comments from the LLM (“here is a 200 word summary; would you like me to make any changes?”; “the data in this table is illustrative, fill it in with the real numbers from your experiments”).*”.

A larger proportion of arXiv papers now contain an AI disclaimer. They must ideally explain what kind of help was generated by these tools (grammar checking, code

development, fact-checking content, develop parts of the paper) and describe the tools and versions used.

Here are some examples using different tones:

- **Karl Svozil** in May 2026: “*The author used an OpenAI Codex language-model-based assistant for structural editing, consistency checks, and implementation of the finite partition census, and accepts responsibility for the final scientific content*” [322].
- **Avantika Agarwal et al.** in May 2026: “*LLMs were used to help write parts of the code to produce numerical results*” [323].
- **Marco Tomamichel** in July 2026: “*Use of artificial intelligence. We are living in strange times. I do not claim to be the first author of this paper; rather, what best describes the situation is that I acted as a busy supervisor giving rough directions to an interesting problem to a student who is brighter and much, much faster than me. A student who then came up with the right proof ideas and fully executed them, wrote the paper and revised it under my guidance. I did not touch any of the writing directly (except for this disclaimer and the acknowledgements) but had the student implement changes according to my feedback, leaving this manuscript in their unique style — as I would usually do. I convinced myself to the best of my ability that all the statements and proofs are correct and checked all the references. The student, in this case, is Claude by Anthropic. The result is interesting enough to deserve wider attention (in my humble judgement), but it is unclear if the standard ways of disseminating work are still appropriate when so much of it is done by artificial intelligence. We will have to decide this as a community; for now, I defaulted to treat this as if it were a normal paper*” [324]. He then commented on X: “*The are two kinds of theory researchers right now: those who have recently played with frontier models and those who still think their jobs will continue to exist in approximately the current form*”.
- **Dominik Peters** in July 2026: “*The reduction was found in July 2026 by GPT 5.6 Sol Ultra and simplified in part with help from Claude Fable 5*” [325].
- **Giulio Malavolta** in July 2026: “*The main idea for this proof was generated biologically, but ChatGPT 5.5 and 5.6 assisted on the proofs of all technical statements involved in the reduction. The final version of this manuscript is written by the author, who takes full responsibility for its content*” [326].
- **Barry Sanders** in August 2026: “*Assistance with manuscript development and editing was provided by ChatGPT (OpenAI GPT-5.6 Sol). Responsibility for all scientific content rests entirely with the author, who independently verified every mathematical and scientific statement*”. [327].

- **Nathan X. Kodama** et al. in August 2026: “*During the preparation of this work, the authors used Claude Opus 4.7, 4.8, and 5 and Claude Code (Anthropic) to improve language and readability, and as coding assistants in developing the research software. All AI-assisted code was written under the authors’ direction and reviewed and tested by the authors before use in any reported experiment. Large language models were not used to design the study, derive the theory, or interpret the results. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the published article*” [328].
- **Andrew M. Childs** in September 2026: “*We (the author and the AI systems that did the heavy lifting) show that the quantum query complexity of minimizing a convex function [...] **AI statement** The results of this paper were generated primarily through conversations with Claude Opus 4.8 and 5, which provided the majority of the technical approach. [...] Claude was then prompted to numerically check this conjecture. [...] Further interactions with Claude Opus 5 and ChatGPT 5.6-Sol helped to illuminate discretization of the query model and the connection to convex optimization. The manuscript was written manually, informed by AI-generated notes. The paper was also checked for correctness and clarity using ChatGPT 5.6-Sol and Claude Opus 5; some technical issues were found, and the paper was revised accordingly. The author takes sole responsibility for correctness of the final manuscript.*” [213].

Peer-reviewing is also being affected by the LLM craze. It is a very demanding task and it becomes sometimes impossible for some broad-spectrum papers. Peer-reviewers are then starting to do their work with the help from LLMs. Like for the papers themselves, you can wonder about the reviewing value if it is not entirely understood by the reviewer. But LLM-based reviewing can be done on one’s own papers to improve them, even in a retrospective way as Maciej Malinowski did by creating an agent in June 2026 with Claude Fable, criss-crossing theory and experimental data reporting [329]. It found errors and he concludes that “*every paper can be improved through feedback from frontier AI models*”!

Just remember one key thing in life in general, and here as well: intelligence starts with the questions you ask, and only after, in the responses you obtain or generate [330–332].

2.6.9 Academic roles

In most countries and in all disciplines, several roles can be distinguished in research organizations.

Doctoral students are students who are undertaking a doctoral thesis (PhD, for Doctor of Philosophy, in latin Philosophiae Doctor, for any science). It lasts from

three to five years depending on the country. This thesis completes a higher education program in the University.

Post-docs or post-doctoral researchers are researchers who, after having obtained their PhD, conduct research in a laboratory under a fixed-term contract. They sometimes do several post-docs in different locations, frequently out of their originating country. It is the anteroom of a full-time research position.

Researchers have a permanent tenure in a research organization whether in the industry or with government funded research organizations. In many countries, they are also civil servant researchers recruited through some open competitions process.

Habilitation to Direct Research (HDR in France) allows a tenured researcher to direct the thesis of one or more doctoral students as a thesis director and to obtain a university professorship. The rules vary from country to country, such as having completed two doctoral theses and having published internationally recognized work in one’s field. The doctorat d’État was abolished in France by the Savary law of January 1984 and this habilitation was created by a decree in November 1988. The HDR is considered to be a diploma. It is awarded on free application by the research commission of the Universities which deliberates in the form of a jury.

Research Directors are researchers with the possibility to autonomously determine the field of their research work. They supervise several doctoral students and post-docs when they are successful with finding the related public and/or private funding. They are also selected by competition in research institutions. Depending on the country and research organization, there are several grades in the function, linked to advancement over time and merit.

Principal Investigators are lead researchers who are in charge of the preparation, conduct, resources allocation and administration of a research grant for which they are the project lead researcher and main holder. Sometimes, a PI is synonym of laboratory director or research group leader.

In addition to these roles, let’s not forget the **laboratory technicians** who set up the experiments and about whom less is said and the **engineers** who can play a role in the creation of many scientific instruments.

2.6.10 h-index

The h-index, named after its creator Jorge Hirsch in 2005, is an index that quantifies a researcher’s productivity and scientific impact. It is based on the level of citations of his scientific publications in peer-reviewed journals. It is a bit like a PageRank for a website, but a simpler one. It is an integer corresponding to the number of papers h that have each obtained at least h citations in other papers (Figure 79).

The level of h-index can be used as a quantitative data for obtaining a position as a resident researcher (10-12), professor (>18) or member of an academy of science (>45).

As with any composite index like Shanghai ranking list of universities, it generates side effects: a race to “publish or perish” papers of little incremental value, cross-referencing between researchers, self-citation, an abundance of co-authors, etc.

The discrepancy of h-index is quite high with researchers with a Nobel prize in physics with low index like with John Clauser (31, Nobel in 2022) and Brian Josephson (24, Nobel in 1973) and very high index like Anton Zeilinger (151, Nobel in 2022), David Wineland (134, Nobel in 2012) or Mikhail Lukin (213), as of September 2026.




Figure 79: h-index explained graphically. An author reaches an h-index of N when N papers have been cited at least N times in other papers. This is not accounting for tricks like self-citations or being secondary authors of papers.

Some alternatives indexes have been proposed like the recent h-frac, but not yet adopted [333]. The correlation between h-index and scientific awards in physics is declining. Thus the proposal for an alternative index named h-frac, for h-fractional, that improves the correlation between the index and other scientometric measures like scientific awards. It allocates citations fractionally and evenly among all coauthors of scanned papers to avoid the phenomenon of low-contribution hyperauthors.

h-index remains, however, an interesting indicator of the influence of researchers and their production volume. On average, the h-index of a researcher in physics is close to the length of his career since his PhD. It obviously evolves over time. It is full of flaws like all quantitative indicators. For example, the basic h-index does not distinguish between the main author and the co-authors. Hence the abundance of authors cited in many papers, some of them having made only marginal contributions.

The index is usually calculated from **Google Scholar** data, but it is sometimes found calculated only on the **Semantic Scholar** website. The most serious index is provided by the **Web of Science** because its database is the cleanest.

2.6.11 Poster sessions

In a scientific conference, a “poster session” is usually a part of the conference dedicated to the presentation of researchers’ projects during a break, in a dedicated area.

Researchers display a poster describing their research work and talk with conference participants as they stroll through the conference exhibition area during dedicated breaks. It is an exercise in humility reminding what Jehovah’s witnesses are doing in the streets.

2.6.12 Figures of merit

This common expression broadly describes a set of specifications and the success metrics to be achieved to bring a given technology to fruition. DiVincenzo’s qubit technology criteria can be considered a figure of merit for success for quantum computing. It usually provides a roadmap and set of goals for researchers and technology vendors.

2.6.13 International collaboration

Nowadays, all modern countries have crafted their “quantum national plan” with a certain willingness to better control their sovereignty. It is like being the first with the atomic bomb during World War II.

But let’s remember that international collaboration between researchers is intense. Most of those I met in French laboratories collaborate with colleagues either in Europe within the framework of Horizon 2020 and now Horizon Europe projects, the European Quantum Flagship or for some ERC grants. They also collaborate with researchers outside the European Union, particularly in Asia (Japan, Singapore), as well as in the USA, UK, Switzerland and Australia.

Quantum science knowledge is quite open and is rather well shared on a global scale. This is encouraged by many international scientific conferences where knowledge is being built, researchers get to know each other, and joint projects are being launched.

This is one of the reasons why I don’t believe in the existence of a supposed quantum computer whose capabilities would defy understanding and which would be hidden in the basement of a secret NSA datacenter to break all the RSA keys of the Internet.

Scientific nationalism in quantum technologies finally comes into play further downstream of research, when it comes to transforming it into industrial advantage.

Technologies often have their “magic sauce”, as in semiconductor manufacturing processes. This has always been the case in digital technologies.

2.6.14 Technology Readiness Level

This technology readiness level notion is commonly used in deep techs. It describes the level of maturity of a technology with a scale from 1 to 9 (Figure 80). It follows a relatively standardized classification whose seven-level ancestor was created by Stan Sadin at NASA in 1974, codified in 1989, extended to nine levels in 1990 and given its current definitions by John Mankins in a 1995 white paper [334, 335]. It was then used by the European Union and various other organizations, and standardized in ISO 16290 for space hardware. It was initially mainly used in the aerospace, defense and energy industries.

This scale can have several use cases. It is used to assess the level of risk and maturity for an investor in a startup. Very advanced deep techs are also the playground of TRL, and quantum technologies are no exception.

The TRL scale has 9 levels [337]:

- **TRL 1:** basic principles are described or observed, at the theoretical or experimental stage.
- **TRL 2:** technological concepts are formulated and not yet necessarily tested.
- **TRL 3:** proof of concept is carried out in a laboratory, at the level of the technical process.
- **TRL 4:** components are validated together in the laboratory.
- **TRL 5:** the technology is validated in a relevant environment, meaning one that is representative of the intended operating conditions.
- **TRL 6:** a technology prototype is demonstrated in this relevant environment.
- **TRL 7:** a prototype is evaluated in an operational environment.
- **TRL 8:** a complete system has been evaluated and qualified.
- **TRL 9:** a complete system is operational and qualified in production.

The relevance of the solution to market needs is missing at this scale, but it is a marketing rather than a technical consideration [338]. Most of the time, it more or less coincides with TRL levels 7 to 9 since reaching this scale requires funding and finding customers willing to test the solution.

Kristel Michielsen has proposed a scale suitable for quantum computing, the QTRL, for the Quantum Technology Readiness Level in Figure 81. Her assessment of

some technologies can be argued. For example, she positions D-Wave’s quantum-annealed computers in TRL 8 and 9. This is commercially correct since these computers are well marketed. This being said, if they are well available physically, it is not proven that they are of much use at the moment.






Figure 81: The quantum TRL scale, created by Kristel Michielsen in 2017. Adapted from [339] by Olivier Ezratty, September 2026.

The TRL scale can be extended with two additional market development steps as shown in Figure 82.

The specificity of quantum technologies is that many hardware startups are created with very low TRLs and are still doing fundamental and applied research in parallel with systems engineering, technology developments and industrialization. This is particularly true for those who are starting to design qubits using technologies that have not yet been proven, even in the labs. Many startups are created by researchers who view the startup vehicle as a relevant one to fund their research. It particularly makes sense when the technology development requires building a team with a critical mass and broad set of skills that can’t be assembled within an academic laboratory team.

In quantum technologies, the notions of “MVP” (minimum viable product) are very different from the classical digital world. It is based on scientific rather than functional metrics. We have many such startups around in quantum technologies because of the famous FOMO (fear of missing out) syndrome with investors.

This shows up with investors who fear missing the future golden goose or unicorn. They are ready to overinvest in companies they perceived will be the future market champion. This explains for example the level of funding for startups like **Rigetti** and **PsiQuantum** or the SPAC funding mechanism (special purpose acquisition company) implemented by **IonQ**, **Rigetti**, **D-Wave** and **Zapata AI**, and the quantum business spin-off from **Honeywell** and its merger with **CQC** which became **Quantinuum** in December 2021.



- [203] [Bridging the NISQ and Fault-Tolerant Regimes: Generative-ML-Assisted Quantum Selected CI for Molecular Simulations](#), by Anurag K. S. V., arXiv, June 2026.
- [204] [Dark matter searches with a 13 meV threshold superconducting sensor array](#), by Christopher Albert, Alejandro Simon, et al., arXiv, July 2026.
- [205] [Withdrawal: A Large-Scale Dataset for Retraction Study](#), by Delip Rao, Jonathan Young, Thomas Dietterich, and Chris Callison-Burch, arXiv, December 2024 (11 pages).
- [206] [Cryptography: Classical versus Post-Quantum](#), by Abhinav Awasthi and Atul Chaturvedi, arXiv, February 2024 (3 pages).
- [207] [Efficient Protein Ground State Energy Computation via Fragmentation and Reassembly](#), by Laia Coronas Sala and Parfait Atchade-Adelemou, arXiv, May 2025 (7 pages).
- [208] [How to write an excellent Review article](#), by Nature editorial board, Nature Reviews Bioengineering, November 2024 .
- [209] [Training of quantum circuits on a hybrid quantum computer](#), by Daiwei Zhu, Christopher Monroe, et al., Science Advances, October 2019 (7 pages).
- [210] [Search for a massless dark photon in \$\Lambda_c^+ \rightarrow p\gamma'\$ decay](#), by BESIII Collaboration: M. Ablikim, M. N. Achasov, et al., arXiv, November 2022 (9 pages).
- [211] [Combined Measurement of the Higgs Boson Mass in \$pp\$ Collisions at \$\sqrt{s} = 7\$ and 8 TeV with the ATLAS and CMS Experiments](#), by G. Aad et al., PRL, May 2015 (35 pages).
- [212] [A method of an on-demand beamsplitter for trapped-ion quantum computers](#), by Takanori Nishi, arXiv, August 2025 (11 pages).
- [213] [Quantum lower bounds for convex optimization and real matrix-vector query problems](#), by Andrew M. Childs, arXiv, September 2026 (29 pages).
- [214] [Fidelity and Entanglement of Random Bipartite Pure States: Insights and Applications](#), by George Biswas, Anindya Biswas, et al., arXiv, May 2024 (6 pages).
- [215] [High-kinetic inductance NbN films for high-quality compact superconducting resonators](#), by Simone Frasca, Edoardo Charbon, et al., arXiv, February 2023 (12 pages).
- [216] [Experimental analysis of energy transfers between a quantum emitter and light fields](#), by Ilse Maillette de Buy Wenniger, Alexia Auffèves, Pascale Senellart, et al., arXiv, June 2023 (18 pages).
- [217] [Dark Matter Fluid Interpretation of Quantum Entanglement](#), by Zhi Cheng, March 2023 (8 pages).
- [218] [The Ecology of Fringe Science and its Bearing on Policy](#), by HM Collins, A Bartlett, and LI Reyes-Galindo, arXiv, June 2016 (27 pages).
- [219] [The entities enabling scientific fraud at scale are large, resilient, and growing rapidly](#), by Reese A. K. Richardson et al., PNAS, August 2025 (11 pages).
- [220] [The fight against fake-paper factories that churn out sham science](#), by Holly Else and Richard Van Noorden, Nature, March 2021 (4 pages) .
- [221] [Retraction Note: Epitaxy of advanced nanowire quantum devices](#), by Sasa Gazibegovic, Diana Car, Hao Zhang, Leo P. Kouwenhoven, and Erik P. A. M. Bakkers, Nature, April 2022 (1 page).
- [222] [Retraction Note: Room-temperature superconductivity in a carbonaceous sulfur hydride](#), by Elliot Snider, Ranga P. Dias, et al., Nature, September 2022.
- [223] [Allegations of Scientific Misconduct Mount as Physicist Makes His Biggest Claim Yet](#), by Dan Garisto, Physics, March 2023.
- [224] [Breakthrough or bust? Claim of room-temperature superconductivity draws fire](#), by Robert F. Service, Science, October 2021 .
- [225] [Superconductivity researcher who committed misconduct exits university](#), by Dan Garisto, Nature, November 2024 .
- [226] [Where are we heading with NISQ?](#), by Olivier Ezratty, arXiv, June 2023 (47 pages).
- [227] [Is there a Moore's law for quantum computing?](#), by Olivier Ezratty, arXiv, March 2023 (32 pages).
- [228] [Mitigating the quantum hype](#), by Olivier Ezratty, arXiv, February 2022 (26 pages).
- [229] [Perspective on superconducting qubit quantum computing](#), by Olivier Ezratty, The European physical Journal-A, May 2023 (18 pages) .
- [230] [Quick thoughts on research](#), by Michael Nielsen, July 2023.
- [231] [COPE Publication Ethics](#).
- [232] Typst is a fast alternative to LaTeX, with near-instant compilation times, simple Markdown-like syntax, clear error messages, and powerful native scripting layout features. It currently faces significant limitations for academic use due to its limited publisher acceptance, a much smaller package ecosystem compared to LaTeX's decades of specialized tools, and a lack of high-end micro-typography controls. It is also instable pre-1.0 software.
- [233] [Artificial Intelligence for Science in Quantum, Atomistic, and Continuum Systems](#), by Xuan Zhang, Shuiwang Ji, et al., arXiv, June 2025 (287 pages).
- [234] [On Hodge-Riemann Cohomology Classes](#), by Julius Ross and Matei Toma, arXiv, June 2021 (23 pages).
- [235] [Dissipative mean-field theory of IBM utility experiment](#), by Emanuele G. Dalla Torre and Mor M. Roses, arXiv, February 2025 (4 pages).
- [236] [Phase transition in Random Circuit Sampling](#), by A. Morvan et al., arXiv, April 2023 (43 pages).
- [237] [A Short Review on Novel Approaches for Maximum Clique Problem: from Classical algorithms to Graph Neural Networks and Quantum algorithms](#), by Raffaele Marino, Lorenzo Buffoni, and Bogdan Zavalnij, arXiv, March 2024 (24 pages).
- [238] [Summary/abstract writing](#), by Nature, 2023 (1 page).
- [239] [Information for Authors](#), by PRX Quantum, PRX Quantum.
- [240] [Why I stopped submitting my work to for-profit publishers](#), by Jonte Hance, PhysicsWorld, May 2025.
- [241] [Quantum Enhanced Inference in Markov Logic Networks](#), by Peter Wittek and Christian Gogolin, Scientific Reports, April 2017.
- [242] [RETRACTION: Origin of the distinct site occupations of H atom in hcp Ti and Zr/Hf](#), by (Fan-Xi Yang et al), ScienceDirect, November 2024 (9 pages).
- [243] [Effects of Research Paper Promotion via ArXiv and X](#), by Chhandak Bagchi, Eric Malmi, and Przemyslaw Grabowicz, arXiv, June 2024 (16 pages).
- [244] [Scientists take step towards quantum supremacy](#), by National University of Science et al., March 2021.
- [245] [Information for reviewers](#), by The Royal Society Publishing - Philosophical Transactions A, 2026.
- [246] [Scientists Take Step Towards Quantum Supremacy](#), by MISIS, March 2021.
- [247] [Quantum sensors for microscopic tunneling systems](#), by Alexander Bilmes, Jürgen Lisenfeld, et al., npj Quantum Information, February 2021.
- [248] [Inside a "Fake" Conference: A Journey Into Predatory Science](#), by Ruairi J Mackenzie, July 2019.
- [249] [LSQuant: Novel Initiative Created To Improve Quantum Transport Methodologies](#), by James Dargan, The Quantum Insider, May 2021.
- [250] [The Franke-Gorini-Kossakowski-Lindblad-Sudarshan \(FGKLS\) Equation for Two-Dimensional Systems](#), by Alexander A. Andrianov, Oleg O. Novikov, et al., arXiv, April 2022 (27 pages).

- [251] Floquet integrability and long-range entanglement generation in the one-dimensional quantum Potts model, by A.I. Lotkov, V. Gritsev, A.K. Fedorov, and D.V. Kurllov, arXiv, April 2022 (10 pages).
- [252] Probing Lorentz-Invariance-Violation Induced Nonthermal Unruh Effect in Quasi-Two-Dimensional Dipolar Condensates, by Zehua Tian, Jiangfeng Du, et al., arXiv, September 2022 (4 pages).
- [253] Emergent quantum mechanics of the event-universe, quantization of events via Denrographic Hologram Theory, by Oded Shor, Felix Benninger, and Andrei Khrennikov, arXiv, March 2023 (18 pages).
- [254] Emergent Sasaki-Einstein geometry and AdS/CFT, by Robert J. Berman, Tristan C. Collins, and Daniel Persson, Nature Communications, January 2022.
- [255] Exploring uberholography, by Dmitry S. Ageev, arXiv, September 2022 (13 pages).
- [256] Variational quantum algorithm for measurement extraction from the Navier-Stokes, Einstein, Maxwell, B-type, Lin-Tsien, Camassa-Holm, DSW, H-S, KdV-B, non-homogeneous KdV, generalized KdV, KdV, translational KdV, sKdV, B-L and Airy equations, by Pete Rigas, arXiv, August 2024 (187 pages).
- [257] Quantum Pontryagin Neural Networks in Gamkrelidze Form Subjected to the Purity of Quantum Channels, by Nahid Binandeh Dehaghani, A. Pedro Aguiar, and Rafal Wisniewski, arXiv, June 2023 (6 pages).
- [258] Multi-self-loop Lackadaisical Quantum Walk with Partial Phase Inversion, by Luciano S. de Souza, Tiago A. E. Ferreira, et al., arXiv, November 2024 (19 pages).
- [259] Efficient and scalable inter-module switching for distributed quantum computing architectures, by Kamil Bradler, arXiv, August 2025 (39 pages).
- [260] Exponential Quantum Speedup for the Traveling Salesman Problem, by Anant Sharma, April 2024 (6 pages).
- [261] Transgressing the Boundaries: Towards a Transformative Hermeneutics of Quantum Gravity, by Alan D. Sokal, Social Text, May 1996 (39 pages) 🗑️.
- [262] Spontaneous Human Combustion rules out all standard candidates for Dark Matter, by Frederic V. Hessman and J. Craig Wheeler, arXiv, April 2023 (5 pages).
- [263] Do Papers with Titles Ending in a Question Mark Usually Have the Answer “No”?, by Daniel Stern and Brian Grefenstette, arXiv, March 2026 (2 pages).
- [264] Terry vs an AI, Round 1: Heralding single-rail (approximate?) 4-GHZ state from squeezed sources, by Terry Rudolph, arXiv, March 2023 (4 pages).
- [265] Computer-inspired Quantum Experiments, by Mario Krenn, Manuel Erhard, and Anton Zeilinger, arXiv, February 2020 (17 pages).
- [266] Digital Discovery of 100 diverse Quantum Experiments with PyTheus, by Carlos Ruiz-Gonzalez, Mario Krenn, et al., arXiv, December 2023 (44 pages).
- [267] Scientists use generative AI to answer complex questions in physics, by Adam Zewe, May 2024.
- [268] Unleashing the Potential of LLMs for Quantum Computing: A Study in Quantum Architecture Design, by Zhiding Liang, Yiyu Shi, et al., arXiv, July 2023 (9 pages).
- [269] Meta-Designing Quantum Experiments with Language Models, by Sören Arlt, Mario Krenn, et al., arXiv, June 2024 (3 pages).
- [270] RydbergGPT, by David Fitzek, Roger G. Melko, et al., arXiv, May 2024 (9 pages).
- [271] Can Theoretical Physics Research Benefit from Language Agents?, by Sirui Lu, J. Ignacio Cirac, Bernhard Schölkopf, et al., arXiv, June 2025 (9 pages).
- [272] A Survey on Large Language Models with some Insights on their Capabilities and Limitations, by Andrea Matarazzo and Riccardo Torlone, arXiv, February 2025 (174 pages).
- [273] CURIE: Evaluating LLMs On Multitask Scientific Long Context Understanding and Reasoning, by Hao Cui, Subhashini Venugopalan, et al., arXiv, May 2025 (48 pages).
- [274] The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity, by Parshin Shojaee et al., Apple Machine Learning Research, June 2025.
- [275] Comment on The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity | alphaXiv, by C. Opus et al., alphaXiv, June 2025 (4 pages).
- [276] Finite entropy sums in quantum field theory, by Mark Van Raamsdonk, arXiv, August 2025 (29 pages).
- [277] Is ChatGPT Transforming Academics’ Writing Style?, by Mingmeng Geng and Roberto Trotta, arXiv, November 2024 (19 pages).
- [278] Threads of Subtlety: Detecting Machine-Generated Texts Through Discourse Motifs, by Zae Myung Kim, Dongyeop Kang, et al., arXiv, June 2024 (26 pages).
- [279] ChatGPT’s explanation of quantum computing is bullsh*t, by Chris Ferrie, January 2023.
- [280] GPT-4 gets a B on my quantum computing final exam!, by Scott Aaronson, April 2023.
- [281] AI and science: what 1,600 researchers think, by Richard Van Noorden and Jeffrey M. Perkel, Nature, September 2023.
- [282] ChatGPT use shows that the grant-application system is broken, by Juan Manuel Parrilla, Nature, October 2023.
- [283] The Economic Climate in Bangladesh, by John E. Spillan et al., Pubpeer, November 2023.
- [284] Analysis of student responses toward activities using e-learning in higher education, by Faiza Rini et al., Pubpeer, May 2023.
- [285] Papers and peer reviews with evidence of ChatGPT writing, by Retraction Watch, March 2024.
- [286] The effect of ChatGPT on students’ learning performance, learning perception, and higher-order thinking: insights from a meta-analysis, by Jin Wang and Wenxiang Fan, Humanities and Social Sciences Communications, May 2025.
- [287] Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task, by Nataliya Kosmyna, Pattie Maes, et al., arXiv, December 2025 (216 pages).
- [288] LLM hallucinations in the wild: Large-scale evidence from non-existent citations, by Zhenyue Zhao, Yihe Wang, Toby Stuart, Mathijs De Vaan, Paul Ginsparg, and Yian Yin, arXiv, May 2026 (16 pages).
- [289] Delving into ChatGPT usage in academic writing through excess vocabulary, by Dmitry Kobak, Jan Lause, et al., arXiv, June 2024 (13 pages) The paper title containing the word “delving”, you can wonder to what extend it was itself written with AI.
- [290] Empirical evidence of Large Language Model’s influence on human spoken communication, by Hiromu Yakura, Iyad Rahwan, et al., arXiv, July 2026 (65 pages).
- [291] AI tools are spotting errors in research papers: inside a growing movement, by Elizabeth Gibney, Nature, March 2025 🗑️.
- [292] Can LLM feedback enhance review quality? A randomized study of 20K reviews at ICLR 2025, by Nitya Thakkar, James Zou, et al., arXiv, April 2025 (30 pages).
- [293] How can AI, LLMs and quantum science empower each other?, by Olivier Ezratty, January 2024 (94 pages).
- [294] Reasoning Language Models: A Blueprint, by Maciej Besta, Torsten Hoeffler, et al., arXiv, June 2025 (44 pages).
- [295] Agentic Reasoning for Large Language Models, by Tianxin Wei, Jingrui He, et al., arXiv, January 2026 (135 pages).
- [296] QuantumQA: Enhancing Scientific Reasoning via Physics-Consistent Dataset and Verification-Aware Reinforcement

- Learning, by Songxin Qu, Zhao-Yun Chen, et al., arXiv, April 2026 (25 pages).
- [297] [Probing the Critical Point \(CritPt\) of AI Reasoning: a Frontier Physics Research Benchmark](#), by Minhui Zhu, Hao Peng, et al., arXiv, May 2026 (40 pages).
- [298] [Reasoning in Action: MCTS-Driven Knowledge Retrieval for Large Language Models](#), by Shuqi Liu, Bowei He, Chen Ma, and Linqi Song, arXiv, November 2025 (10 pages).
- [299] [DeepScientist: Advancing Frontier-Pushing Scientific Findings Progressively](#), by Yixuan Weng, Minjun Zhu, Qiujie Xie, Qiyao Sun, Zhen Lin, Sifan Liu, and Yue Zhang, arXiv, September 2025 (19 pages).
- [300] [Evaluating Large Language Models in Scientific Discovery](#), by Zhangde Song, Chenru Duan, et al., arXiv, May 2026 (63 pages).
- [301] [Accelerating Scientific Research with Gemini: Case Studies and Common Techniques](#), by David P. Woodruff, Vahab Mirrokni, et al., arXiv, March 2026 (154 pages).
- [302] [Quantum-Audit: Evaluating the Reasoning Limits of LLMs on Quantum Computing](#), by Mohamed Afane, Juntao Chen, et al., arXiv, February 2026 (18 pages).
- [303] [Why AI systems don't learn and what to do about it: Lessons on autonomous learning from cognitive science](#), by Emmanuel Dupoux, Yann LeCun, and Jitendra Malik, arXiv, March 2026 (34 pages).
- [304] [Why Do Multi-Agent LLM Systems Fail?](#), by Mert Cemri, arXiv.
- [305] [The Illusion of Insight in Reasoning Models](#), by Liv G. d'Aliberti and Manoel Horta Ribeiro, arXiv, April 2026 (43 pages).
- [306] [Products of consecutive integers with unusual anatomy](#), by Terence Tao, arXiv, April 2026 (41 pages).
- [307] [Mathematics in the age of AI](#), by Terence Tao, July 2026 (52 slides).
- [308] [Mathematics in the age of AI](#), by Terence Tao, arXiv, August 2026 (12 pages).
- [309] [Ten Advances in Mathematics and Theoretical Computer Science](#), by OpenAI, July 2026 (253 pages).
- [310] [A curated list of 219 awesome papers exploring the use of artificial intelligence / machine learning / deep learning for mathematical discoveries](#), GitHub, August 2026.
- [311] [Auditing an AI-Generated Mathematical Proof: A Correction to a Greedy Conditioning Lemma in Quantum Parallel Repetition](#), by Mikołaj Sienicki and Krzysztof Sienicki, arXiv, August 2026 (8 pages).
- [312] [The crisis of AI-generated mathematics](#), by Max Weinreich, arXiv, August 2026 (7 pages).
- [313] [Early science acceleration experiments with GPT-5](#), by Sébastien Bubeck, Nikita Zhivotovskiy, et al., arXiv, November 2025 (89 pages).
- [314] [Lowering the implementation barrier of neutral-atom quantum computing with agentic workflows](#), by Constantin Dalyac, Alexandre Dauphin, Loïc Henriet, and Christophe Jurczak, arXiv, July 2026 (10 pages).
- [315] [Evolutionary Discovery of Bivariate Bicycle Codes with LLM-Guided Search](#), by Juan Cruz-Benito, Andrew W. Cross, David Kremer, and Ismael Faro, arXiv, June 2026 (42 pages).
- [316] [Machina Mirabilis](#), by Michael Hla, Michael Hla blog, March 2026.
- [317] [DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning](#), by DeepSeek-AI and Daya Guo et al, arXiv, January 2026.
- [318] [LLMs as models for analogical reasoning](#), by Sam Musker et al., Journal of Memory and Language, December 2025 (19 pages).
- [319] [Inductive or Deductive? Rethinking the Fundamental Reasoning Abilities of LLMs](#), by Kewei Cheng, Yizhou Sun, et al., arXiv, August 2024.
- [320] [Wiring the 'Why': A Unified Taxonomy and Survey of Abductive Reasoning in LLMs](#), by Moein Salimi, Shaygan Adim, Danial Parnian, Nima Alighardashi, Mahdi Jafari Siavoshani, and Mohammad Hossein Rohban, arXiv, April 2026 (48 pages).
- [321] [ArXiv bans AI-generated papers: How science adapts to AI revolution](#), by Jean Piot, Les e-novateurs, August 2026.
- [322] [Answer Partitions and Oracle Access Determine Quantum Query Complexity](#), by Karl Svozil, arXiv, July 2026 (5 pages).
- [323] [Enhanced quantum capacity thresholds from symmetry](#), by Avantika Agarwal, Amolak Ratan Kalra, Sungjai Lee, Debbie Leung, Luke Schaeffer, Pulkit Sinha, and Graeme Smith, arXiv, May 2026 (25 pages).
- [324] [A strong converse for stabilizer codes over Pauli channels via the blowing-up lemma](#), by Marco Tomamichel, arXiv, July 2026 (26 pages).
- [325] [Kemeny Rank Aggregation is NP-Hard for Three Voters](#), by Dominik Peters, arXiv, July 2026 (15 pages).
- [326] [Quantum Separability in Polynomial Time](#), by Giulio Malavolta, arXiv, July 2026 (24 pages).
- [327] [The Dirac Information Carrier for Relativistic Quantum Computation](#), by Barry C Sanders, arXiv, August 2026 (6 pages).
- [328] [Scalable quantum simulation of continuous-time generative models via tensor networks](#), by Nathan X. Kodama, L. Andrew Wray, Sam Cochran, Chad Rigetti, Shravan Veerapaneni, and Michael J. Keiser, arXiv, August 2026 (33 pages).
- [329] [AI for peer review, ft. Claude Fable](#), by Maciej Malinowski, Reading the quantum, June 2026.
- [330] [Towards an AI co-scientist](#), by Juraj Gottweis, Vivek Natarajan, et al., arXiv, February 2025 (81 pages).
- [331] [Lean Copilot: Large Language Models as Copilots for Theorem Proving in Lean](#), by Peiyang Song, Kaiyu Yang, and Anima Anandkumar, arXiv, May 2025 (26 pages).
- [332] [Can ChatGPT be used to generate scientific hypotheses?](#), by Yang Jeong Park, Ju Li, et al., arXiv, March 2023 (37 pages).
- [333] [The h-index is no longer an effective correlate of scientific reputation](#), by Vladlen Koltun and David Hafner, arXiv, February 2021 (26 pages).
- [334] [Technology Readiness Levels at 40: A Study of State-of-the-Art Use, Challenges, and Opportunities](#), by Alison Olechowski, Steven D. Eppinger, and Nitin Joglekar, PICMET Proceedings, August 2015 (11 pages) 📄.
- [335] [In search of technology readiness level \(TRL\) 10](#), by Jeremy Straub, Aerospace Science and Technology, November 2015 (9 pages) 📄.
- [336] [Some explanations on the TRL \(Technology readiness level\) scale](#), by DGA, 2009 (15 pages).
- [337] [Technology Development Stages and Market Readiness](#), by Surya Raghu, June 2017 (35 slides).
- [338] [TRL, MRL, POC, WTF?](#), by Massis Sirapian, April 2019.
- [339] [Simulation on / of various types of quantum computers](#), by Kristel Michielsen, December 2017 (40 slides).
- [340] [Ammonia Technology Roadmap – Analysis](#), by IEA, IEA, October 2021 (168 pages).

This **Research for dummies** subsection contains 153 bibliographical references and notes containing at least 3,589 pages.

2.7 Quantum physics history and scientists key takeaways

- A first wave of 19th century scientists laid the groundwork that helped create quantum physics afterwards (Young, Maxwell, Boltzmann, mathematicians). The photoelectric effect, black body spectrum, atomic emission or absorption spectra as well as the Zeeman effect could not be explained with the theoretical frameworks of classical physics and thermodynamics of the 19th century.
- Starting with Max Planck in 1900, a second wave of scientists (Einstein, De Broglie, Schrodinger, Heisenberg, Dirac, Born, Von Neumann, ...) created the field of quantum mechanics, or quantum physics, to describe light-matter interactions, energy quantization and wave-particle duality. It explained many previously unresolved 19th century experiments while raising new questions about the nature of reality at the nanoscale. 1925 was a pivotal year with the work of Werner Heisenberg, Erwin Schrödinger and Max Born, who laid much of the groundwork of quantum physics as it is still used today.
- Many of these scientists made major theoretical contributions, while other researchers made landmark experimental discoveries such as superconductivity, electron interference and the Stern-Gerlach experiment. Some scientists contributed to both theory and experiment. Quantum physics also relies on a significant body of mathematics, including linear algebra and group theory. After World War II, enabling device technologies such as transistors and lasers increasingly relied explicitly on quantum physics, as part of what is now called the first quantum revolution.
- Since the 1980s and thanks to advances in individual quantum objects control (atoms, electrons, photons) and the usage of quantum superposition and entanglement, new breeds of technologies were created, most of them belonging to the “quantum information science” field and being part of the second quantum revolution. Many of these research programs were funded by governments after Peter Shor’s integer factoring algorithm was created.
- This part showcases many contemporary quantum scientists in quantum physics and quantum information science from around the world with their key contributions. This hall of fame covers many generations, including young brilliant scientists in their 30s.
- While the first quantum revolution was driven by research coming mostly out of Europe, the last wave comes out of all countries across several continents (North and South America, Europe, Asia/Pacific).
- This book also describes how research works in general and particularly in quantum physics and information science. It explains how scientific papers are written and communicated, how researchers are evaluated, and how quantum-technology readiness level can be assessed. It provides some insights on how to analyze and rate scientific papers. It also describe how LLM-based generative AI is transforming, if not disrupting, this whole field like most scientific fields.

3 Quantum physics 101

So far in the book, we have looked at quantum physics and quantum technologies through the lenses of History and key scientist contributions. We discovered how the quantum physics formalism could explain past experiments results where classical physics (statistical physics, thermodynamics) couldn't find a clear explanation. We here have a new look at the fundamentals of quantum physics in a more structured way. Several years of undergraduate and graduate studies are usually necessary to master quantum physics notwithstanding its rich mathematical foundations. This part provides some background knowledge that will help you better understand the various quantum information systems and technologies exposed in the remainder of this book. You will discover a lot of new scientific lingua, even if you have already a preliminary knowledge of what is a qubit or a quantum computer.

Quantum physics appeared at the beginning of the 20th century to explain the dynamics of elementary particles, like **photons**, **electrons** and **atoms** (Figure 83). The dimensions of elementary particles are 10^{-10} m for an atom and 10^{-15} m for the diameter of a hydrogen atom nucleus, thus of a single proton. The electron is a dimensionless pointlike particle according to the standard model although its potential size was upper bounded at 10^{-22} m by Hans Dehmelt in 1988. From the Planck scale, where quantum-gravity effects are expected to become important, up to the atomic scale, quantum theory operates across a range of approximately 10^{25} .

Quantum physics also deals with elementary particles from the standard model like quarks and neutrinos, but it is usually out of the scope of the “second quantum revolution” [341].

and orbital contributions. Nucleus spin plays a role in various quantum technologies like with NV center and donor spins qubits. We also care about it with electron spin-based qubits since nucleus spin can have a detrimental impact on electron spins handling qubits information. It relates to the kinds of isotopes of carbon and silicon that are used in carbon nanotubes and silicon wafers used to create electron spin qubits.

Although sometimes mysterious and said to be “incomplete” by some, quantum physics has gone through the test of time and experiments for over a whole century. Thousands of experiments have validated the underlying postulates, theory and mathematical formalism behind it even though we still cannot clearly describe the physical nature of quantum entanglement, of the electron spin and of the wave-particle duality phenomenon.

Quantum physics first helped explain various observations such as the **black-body radiation** (solved by Max Planck in 1900), the **photoelectric effect** (solved by Albert Einstein in 1905) and the **sharp spectral lines** observed with excited atoms like hydrogen (solved by Niels Bohr and its atom model in 1913 and later with the discovery of the electron spin). Later on, in the mid 1920's, quantum physics was built upon a **mathematical formalism** using multi-dimensional Hilbert spaces and vectors. It centered around the **Schrödinger wave equation** which describes how a massive particle like the electron behaves over space and time, using complex number probability amplitudes and differential equations over time and space. These provide a probabilistic insight on the outcome of the measurement of a particle's energy, momentum, and many other physical properties.

Quantum physics differs from classical physics with demonstrating how and why quantum particles energy, momentum, angular momentum and other metrics are restricted to discrete values (**quantization**), objects behave as particles or waves depending on the context (**wave-particle duality**), and there are limits to how accurately the value of a physical quantity can be predicted prior to its measurement, given a complete set of initial conditions (**complementarity**, **conjugate**




Figure 83: What particles are we dealing in quantum physics? All of them, but in the second quantum revolution, we mainly use electrons, photons, atoms and their nucleus spin. Source: Wikipedia.

In some cases, we still care about atom an atom's nucleus spin, which is the total angular momentum from the spins and orbital motion of its nucleons which are neutrons and protons. A nucleon spin has quark, gluon,

variables, indeterminacy principle) (Figure 85 and Figure 84).

It also refers to **state superposition** which is at the basis of qubit operations which enables interferences among computational amplitudes, and **entanglement** which is a consequence of superposition applied to several quantum objects and shows up with composite quantum systems whose joint state is mathematically nonseparable, while many multipartite superpositions are unentangled. Entanglement is used with multi-qubits quantum gates and is also used in quantum communications and cryptography, and even quantum sensing. Quantum objects **no-cloning** is a particular aspect of quantum physics that limits what we can do with qubits and how memory is managed. Finally, **quantum tunnelling effect** has some impact in quantum technologies, like with the Josephson junctions used in superconducting qubits and with D-Wave quantum annealers. We will also see that there is an exception to this definition with **continuous variable quantum systems** which use observables with continuous values, such as field quadratures, within the same quantum formalism. Other observables, including photon number, may be discrete.

Quantum physics explains other physical phenomena belonging to the broad **quantum matter** category which can be of macroscopic scale compared to atomic and electron particle scales, like **superconductivity** which plays a key role in superconducting qubits, **superfluidity**, used with liquid helium in cryostats and **quantum vacuum fluctuation** and its role in quantum decoherence. It also enabled the creation of **lasers**, used everywhere like for controlling cold atom and trapped ion qubits and for all photonic based quantum technologies. **Polaritons** are hybrid light-matter quasiparticles formed through strong coupling between photons and material excitations. It can become useful in quantum sensing and quantum simulations. The quantum objects bestiary also includes **skyrmions** and **magnons**!



Figure 84: Taiwan published a stamp to celebrate the 2025 International Year of Quantum Science and Technology. It contains references to a wealth of scientific discoveries: the black-body radiation and Planck equation, the photo-electric effect, the Bohr atomic model, entanglement, the Bloch sphere, etc. [Source](#).

3.1 Postulates

Quantum physics formalism is based on a set of postulates that follows. Why are these postulates and not laws? Mainly because they describe a mathematical formalism that cannot be proven per se.

One of the other reasons is that quantum physics does not rely on an ontology describing the physical objects it is based upon. Whenever possible, I will try to connect these postulates with some physical meaning. If all of this seems gibberish to you, skip it!

Postulate I - Quantum state: the state of an isolated physical system is represented, at a given time t , by a state vector $|\psi\rangle$ (psi) belonging to a Hilbert space H called the state space with vectors of length 1, using complex numbers. This is the canonical definition of a quantum state. The $|\psi\rangle$ vector contains the knowledge we can have of a quantum system, represented by the values taken by its measurable and compatible properties. A broader definition of a quantum state is the ensemble of values taken by compatible physical properties of a system made of one or several quantum objects. These compatible properties are measurable simultaneously or in any order.

More generally, a quantum state is represented by a positive semidefinite, unit-trace density operator ρ . Mixed states and reduced states of entangled subsystems require this more general description.

The $|\psi\rangle$ vector is a mathematical object that helps determine and predict over time the probabilistic distribution of the various values of the quantum object compatible properties. The immediate consequence of this first postulate is the notion of superposition where a linear combination of several $|\psi\rangle$ vectors can form another valid quantum state. For a generic qubit, its quantum state defines its amplitude and phase as seen later in the Bloch sphere description. $|\psi\rangle$ is then a vector in a two-dimensional Hilbert vector space combining the $|0\rangle$ and $|1\rangle$ basis states with their related complex amplitudes.

Postulate II - Physical quantities: are related in quantum physics with observables that are both a measurable physical quantity such as the momentum of a particle, and the Hermitian operator to represent it mathematically, as $\hat{A}$ acting on the $|\psi\rangle$ vector in $\hat{A}|\psi\rangle$. $\hat{A}$ is a Hermitian linear matrix operator acting on the state vector $|\psi\rangle$ to evaluate quantized or continuous physical properties of quantum objects. It is a self-adjoint matrix with real eigenvalues. Pauli matrices σ_x , σ_y or σ_z are Hermitian unitary observables. Meanwhile, a projector is a mathematical operator P that projects a quantum state onto a particular subspace. An orthogonal projector is both Hermitian ($P = P^\dagger$) and idempotent ($P^2 = P$). Its physical counterpart is that an immediate repetition of an ideal projective measurement yields the same eigenvalue, when there is no additional evolution. According to the spectral




Figure 85: Eight key basic principles (quantization, wave-particle duality, state superposition, entanglement and indetermination), postulate (measurement), effects (tunnel effect) and theorem (no-cloning) of quantum physics. © Olivier Ezratty, 2021-2025.

theorem, an A can be expressed as a linear combination of mutually orthogonal projectors:

$$A = \sum_i a_i P_i$$

A is the observable being measured, a_i are the distinct, real eigenvalues (the actual physical quantities you read), and P_i are the projectors on the eigenspaces associated with each a_i .

Let's define properties and their variations:

Properties correspond to a quantum system's various observables. For a photon, it can be, for example its phase, polarization, and wavelength, relative to the observation (like the polarizer angle, or observer position). While it is theoretically possible to evaluate the values of all quantized properties of quantum systems like the electron to describe it, continuous properties like a quantum object momentum and position can each be measured, it can't be done with an arbitrary precision for both value due to Bohr's complementarity principle.

Exclusive property values are the possible results of a quantum measurement of a quantized property. The classical examples are vertical and horizontal polarization for a photon or spin up or down for an electron spin along a projection axis. These are mutually exclusive since it corresponds to two results of a physical measurement. Mathematically speaking, two properties are exclusive if their projector operators (*aka* observables...) are orthogonal. Otherwise, these are non-exclusive properties.

Compatible properties of a quantum system can be measured in any order or simultaneously [342]. In that case, their observable operators A and B com-

mute ($AB = BA$), or their commutator is equal to zero ($[A, B] = AB - BA = 0$) [343]. Compatible properties have commuting observables. The notion of properties compatibility must not be confused with complementarity. There is complementarity between incompatible properties, like position and momentum! Measuring a complete set of commuting observables (CSCO) constitutes the most complete measurement of a quantum system. One example is the momentum and energy of a free particle.

Incompatible properties *aka* conjugate variables cannot be measured simultaneously and their observable operators A and B do not commute ($AB \neq BA$ or $[A, B] \neq 0$). This is a particularity of quantum mechanics. Typical examples are continuous variables like momentum and position for a particle or in the discrete variable cases, the spin components of a particles along different directions. However, revealing one property value with a measurement does not exclude revealing another property afterwards. But it is not possible to obtain exact knowledge of both properties at the same time (in the probabilistic sense and following Born's rule). For a single particle, after measuring the X spin component, a Z measurement will yield a totally random result. Also, the energy and position of an electron are incompatible continuous properties.

Postulate III - Measurement: is the result of a physical quantity measurement with an observable operator A . The measurement result is one of the observable operator eigenvalues. We define eigenvalues later and cover the related mathematical formalism in the measurement section of this book. This postulate is sometimes embedded or associated with the previous one. The observable operator doesn't generate a measurement result per se. It helps create a probabilistic distribution of the possible measurement outcomes of a property given

what is mathematically known of the quantum object state vector. When applied to a quantum object vector, it creates another state vector along the eigenvectors of the observable operator. Its spectral projectors, together with the state, determine the probabilities of the possible outcomes. The **expectation value**, or **predicted mean value**, is the average value of repeated measurements that would be obtained with the physical implementation of the observable. The measurement postulate is also named the von Neumann measurement postulate.

Postulate IV - Born rule: when the physical quantity A is measured on a quantum system in a normalized pure state $|\psi\rangle$, the probability of obtaining an eigenvalue α_n for discrete values or α for continuous values of the corresponding observable A is given by its squared amplitude modulus of the related wave function. This is related to Max Born's probability rule. A quantum state can be generally represented by a density operator, which is a square matrix, nonnegative self-adjoint operator ρ normalized to be of trace 1. The average expected value of A in the state ρ is $\text{tr}(A\rho)$, the trace (sum of diagonal matrix values) of the observable operator applied to the density matrix. There are variations of this postulate for various quantum spectrum (discrete and nondegenerate, discrete and degenerate, continuous and non-degenerate). This postulate is sometimes merged with the measurement postulate. This postulate is associated with the principle of spectral decomposition. For a single qubit, the Born rule is simple to describe with $|\alpha|^2$ being the probability of getting a $|0\rangle$ and $|\beta|^2$ of getting a $|1\rangle$ when the qubit state is described as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ with α and β being complex numbers. And due to probabilities normalization, $|\alpha|^2 + |\beta|^2 = 1$.

Postulate V - State collapse: only one result is obtained after a quantum measurement. An ideal projective measurement on the computational basis which is repeated immediately, with no intervening quantum evolution, returns the same eigenvalue. For a qubit, after we measure its state, whatever it is, we get a $|0\rangle$ or a $|1\rangle$ and this becomes the new qubit state after measurement, even if you repeat the measurement. But let's say you obtain a $|1\rangle$, it can then drift back to $|0\rangle$ due to spontaneous energy relaxation.

Postulate VI - Time evolution: the time evolution of the state vector $|\psi(t)\rangle$ of a closed system is governed by the Schrödinger wave equation. As a result, the postulates are applicable for massive non-relativistic particles. Relativistic massive particles time evolution is described by the Dirac and Klein-Gordon equations while photons are covered by Maxwell's equations and their various derivations.

We don't directly deal much with time evolutions to understand quantum computing with qubits and gates, but it still plays a key role in quantum annealing and quantum simulation and, behind the scenes, in gate-based computing, with qubits decoherence, quantum

noise, quantum error corrections mechanisms and measurement.

There is also a **Composition postulate**, which defines the notion of tensor product applied to separable composite quantum systems, *aka* "composite systems" with John Preskill's axioms. We will describe it when covering linear algebra and qubit registers in the next part of this book (section 4).

There are many variations of these postulates in shape, form, name and number, which ranges from 4 to 9 depending on the source [344] (Figure 86). Quantum State can become State Space and Physical Quantities become Unitary Dynamics. John Preskill lists five 'axioms', considering that postulates are axioms since they are not contradicted experimentally [345]. There is not really a single "bible" of quantum postulates even when reading quantum physics founders writings (Bohr, Heisenberg and others) who didn't agree on all of it. I have consolidated below a table with some of these variations of postulates. Imagine if there were various versions of the Bible with 5, 7, 9, 10 and 12 commandments!

Mostly covered in the linear algebra section, the main related quantum physics mathematical tools are:

- **Linear algebra** with complex numbers, eigenvectors, eigenvalues and eigenstates (in 4.1).
- **Functional analysis** with Hilbert spaces, Hermitian matrices, linear operators, spectral theory (also in 4.1).
- **Harmonic analysis** with Fourier transforms and series (in 4.1.12).
- **Groups theory**, including Lie groups (in 4.1.13).
- **Differential equations** with partial differential equations, separation of variables, ordinary differential equations, Sturm–Liouville problems and eigenfunctions (we cover these in the quantum software part of this book).

We consolidate an inventory of the key differences between classical physics and quantum physics in Table 2.

3.2 Quantization

In quantum physics, material or field quantum objects have some physical properties that have discrete values, unlike distances and speeds in classical physics. This corresponds for example to electron energy levels for atoms, but also deals with electrons, atom nucleons and nucleus and photons various quantum properties like their angular momentum and spin (Figure 87). The particles from the standard model (quarks, gluons, neutrinos, ...) that are studied in high energy particles physics (HEP) also have their quantized properties, but we won't deal with it here.

quantum physics postulates variations




Figure 86: A compilation of various inconsistent lists of quantum postulates and axioms. All these postulates formulations have some quantum state and systems dynamics postulates. There is some variable consolidation of physical quantities, observable operators, measurement, state collapse and Born rule postulates. Some quantum scientists add various other postulates in light blue like composite systems or fermions asymmetry. Time evolution usually relies on Schrodinger's simplified wave equation with an unspecified Hamiltonian, and works with massive particles and photons. © Olivier Ezratty, 2022-2025.

Criterion	Level	Classical physics	Quantum physics
State description	Plain	A physical object has a definite position and velocity at every instant, which can be predicted relatively to an observer position.	A quantum object's property is undefined until it is observed in a given measurement context, whether being discrete (electron's quantum numbers) or continuous (position, velocity, phase).
	Formal	A point $(\mathbf{q}, \mathbf{p})$ in phase space for a classical harmonic oscillator.	A wave function $\psi(x, t)$ in a Hilbert space, with $ \psi ^2$ being a probability density.
Predictability	Plain	Knowing a present physical property helps describe its past and future value in a deterministic way. E.g.: planets movement around the Sun.	Only the probabilities of measurement outcomes can be predicted, not the outcome itself. E.g.: measuring a spin in a Stern-Gerlach experiment (Figure 44).
	Formal	Unique trajectory set by initial $(\mathbf{q}_0, \mathbf{p}_0)$.	Born rule with probability $P = \langle \phi \psi \rangle ^2$.
Evolution	Plain	Smooth motion driven by forces.	The quantum wave evolves smoothly between measurements.
	Formal	Hamilton equations, $\mathbf{F} = d\mathbf{p}/dt$.	Schrödinger equation for quantum objects having a mass, $i\hbar \partial_t \psi = \hat{H} \psi$ (Figure 108).
Wave-particle duality	Plain	Waves and particles are distinct things, only waves generating interferences.	A quantum object is both wave and particle, generating interferences between them and with themselves, and momentum on other quantum objects.
	Formal	No self-superposition. A cat can't be dead and alive simultaneously.	Superposition principle with wave function amplitudes additions and interferences.
Conjugate variables	Plain	A physical object position and momentum can both be known with no limit.	Position and momentum are incompatible (see 3.5).
	Formal	Poisson bracket $\{q^i, p_j\} = \delta_j^i$.	Commutator $[\hat{q}^i, \hat{p}_j] = i\hbar \delta_j^i$.
Uncertainty	Plain	No built-in floor on joint precision measurement for position and velocity.	A fundamental precision trade-off between position and velocity at the quantum scale.
	Formal	$\Delta q \Delta p$ may vanish.	$\Delta q \Delta p \geq \hbar/2$ (Heisenberg indeterminacy, seen in 3.5).
Continuity vs discretization	Plain	Energy and related quantities vary continuously, like the speed of a car on a road.	Electron energy, momentum and angular momentum come in discrete steps within the context of an atom. It shows up in atoms spectrography like with the hydrogen atom.
	Formal	Continuous values and spectra.	Discrete eigenvalue spectra, e.g. levels E_n .
Classical limit	Plain	Classical physics is what quantum physics becomes for large, heavy, everyday objects.	A quantum system looks classical once its action far exceeds $\hbar$.
	Formal	Hamilton-Jacobi action $S(\mathbf{q}, t)$, the phase of the classical motion.	$\psi \sim e^{iS(\mathbf{q}, t)/\hbar}$, returning the classical path when $S \gg \hbar$ [346].
Typical regime	Plain	Projectiles, planetary orbits, rainbows, real cats.	Atomic energy levels, photon emission, tunneling.
	Formal	Action $S \gg \hbar$, e.g. a 1 kg mass over 1 m in 1 s gives $S \approx 0.5$ J·s, about $5 \times 10^{33} \hbar$.	Action $S \sim \hbar$, e.g. an electron bound in an atom, where S over one orbit is of order $h = 2\pi\hbar$.

Table 2: comparing classical physics versus quantum physics with a plain-language level and a more formal level. In the first row and later, q or x denote an object position and p its momentum. ψ is a measured quantum state, and ϕ is the measurement outcome, which is an eigenstate of ψ for an observable A . In the bottom rows, S is an action in J·s. Indeed, the Planck constant is energy $\times$ time, or momentum $\times$ length, and also angular momentum.




Figure 87: The three kinds of properties for atoms, electrons, nucleons and photons: their intrinsic properties which are related to the particle type and have a single possible value (mass, charge, spin magnitude), the relational quantized properties which are relative to a system (the electron in an atom) or observation (spin projection on an axis, the photon relative to an observer), and the dynamic variables which are continuous and relative to a referential. The relational properties are considered to be “degrees of freedom”. The table shows the correspondence between atoms, electrons, nucleons and photons. Only elementary particles like electrons and photons have “quantum numbers”, atoms and nucleons being composite quantum objects made of electrons and quarks only have “quantum properties”. The properties in the purple zone are those which are quantized. The photon wavelength and phase are continuous variables whose value is relative to the observer, whereas the OAM is quantized, taking integer multiples of $\hbar$. The polarization is a bit subtle. The helicity or circular polarization is observer independent while a linear polarization is dependent on the observer, depending on the chosen polarizer angle (more on photon properties in 3.2.4). Quantized properties are used to create qubits or qutrits in quantum computing with distinct states and at the particle scale (atoms, electrons, photons). © Olivier Ezratty, 2022-2026, with some Wikipedia images sources.

As shown in Figure 87, quantum objects like atoms, nucleons, electrons and photons have two types of properties:

- **Intrinsic structural properties** which define what the particle is. They are independent of the environment or of the observer. They are the object mass and charge which are fixed constants, the spin magnitude, which is the “internal” angular momentum and is fixed per particle type ($s = 1/2$ for the electron and the nucleons, $s = 1$ for the photon). Since they have only one possible value, there is no possible superposition or entanglement using these properties. Mass and charge are considered to be classical properties and not quantum per se. The spin magnitude is intrinsic but considered to be quantum since it has no classical equivalent.
- **Extrinsic relational properties** which define the particle state in relationship to its environment. These states can be superposed until a measurement occurs. However, while a state can be expressed as a superposition in a chosen basis, measurement statistics and the postmeasurement state depend on the measurement basis. These properties are also quantum “degrees of freedom”. There is then a distinction between quantized properties and continuous variables:
 - **Quantized properties** or “discrete quantum variables”, like the electron principal number, its orbital quantum number and orbital angular momentum which characterize its relationship with a nucleus in an atom. These are “extrinsic” because they describe the motion of the electron in the context of the atom. Photons quantized properties are their circular polarization, orbital angular momentum (OAM) and Fock number which counts the number of photons having the same physical properties. These variables are in a finite dimension Hilbert size like, as we will see later, 2^n for n qubits. Such quantum variables are used in gate-based quantum computing and with some quantum key distribution techniques, labeled “discrete variable quantum key distribution” (DV-QKD).
 - **Continuous variables** like the quantum object position and velocity which are relative to a frame of reference (an observer) and in the case of the photon, its wavelength, phase and linear polarization. These variables are in a Hilbert space of infinite dimension. They are used in continuous variables quantum computing like with Xanadu’s qumodes, and with continuous variable quantum key distribution (CV-QKD). How can we have “quantum” objects with properties that are not “quantized” in the first place? It still belongs to the quantum world and its mathematical framework with Hilbert spaces, quantum states and operators. They exhibit wave-particle duality, superposition, entanglement, probabilistic measurement

and are subject to the Heisenberg indeterminacy principle, which is nearly always illustrated with continuous variables like a particle position and momentum. Continuous quantum variables are also quantum in opposition to the macroscopic analog world where all these quantum phenomena like the wave-particle duality are not observed.

Is time a discrete or continuous physical property? The debate is not settled yet. In quantum physics, time is usually not discretized, like in Schrödinger’s wave equation. Still, some experiments try to prove that it can be discretized, of course, at a very low scale and in connection with quantum gravity [347, 348]. Then you have Planck’s time, 5.39×10^{-44} s, often quoted as a lower bound. It is not an experimentally established minimum duration, but the scale at which a quantum theory of gravity is expected to be required.

In quantum information systems, we use quantum objects which can usually have two different separable states that can be initialized, modified, and measured. Superconducting qubits rely on two separable and controllable energy and phase levels for the oscillating current flowing through the insulator in Josephson junctions.

3.2.1 Black body spectrum explanation


There is a correspondence between the discontinuous energetic transitions of electrons in orbit around atoms and the related absorbed or emitted photons. Quantization shows up in other various places like crystals. Atoms also form harmonic oscillators and vibrate at quantified amplitudes in crystalline structures, according to a model Einstein developed in 1907. You will find many quantum oscillators all over the place, like in superconducting qubits and various forms of cavities.

Quantization was a way to progressively explain experiments done beforehand, the first being about the blackbody radiation spectrum (Figure 88). This one marked the beginnings of quantum physics.

Before explaining black body spectrum, let’s recall the three kinds of spectrum that can be usually found experimentally and are pictured in Figure 89.


- A **continuous spectrum** comes from a hot and dense body like the sun, a heated solid or a perfect such body *aka* black body. It contains light in all visible frequencies that come from the random excitement of atoms in the examined body.
- An **absorption spectrum** is usually made of a continuous source of light traversing an absorbing medium like a cold gas. The resulting spectrum will be a continuous one with black lines corresponding to the frequencies absorbed by the medium.
- An **emission spectrum** is created by some rarefied hot gas. It shows discrete spectrum lines correspond-

quantization = discontinued nanoscopic properties of nanoscale matter and light




explained the **black body** electromagnetic radiation spectrum and fixed the **ultraviolet catastrophe** from the classical theory

B. Stewart, 1858 and G. Kirchhoff, 1859
Max Planck, 1900



explained the **photoelectric effect** with UV rays extracting electrons from a metal surface in vacuum

H. Hertz, 1887 and P. Lenard, 1902
Albert Einstein, 1905



explained the emission of **sharp spectral lines** by gas atoms in an electric discharge tube with discrete electron **energy level transitions**

Johan Balmer, 1885
Theodore Lyman, 1906
Friedrich Paschen, 1908
explanation: **Niels Bohr, 1913**

Figure 88: The three fundamental 19th century electromagnetic waves experimental results which were later explained by quantum physics, all explained by quantization of the electromagnetic wave field. Green labels correspond to experiments and black to theory © Olivier Ezratty 2023-2026 compilation. Various sources.




Figure 89: Differences between continuous spectrum, absorption spectrum and emission spectrum. Emission spectrum analysis is one key tool used in astrophysics to understand the composition of stars and even of exoplanets, using for example the transit method, comparing the spectrum of stars before and after they are partially hidden by an exoplanet passing in between the star and the Earth. © Olivier Ezratty, 2021.

ing to photons emitted by the excited gas atoms at specific frequencies.

Black bodies were theorized by **Gustav Kirchhoff** in 1859. These are ideal physical bodies in thermal equilibrium that absorb all incident electromagnetic wave radiations and reflects or transmits none. Since it absorbs all wavelengths, it is supposed to be black, although stars like the Sun are good approximations of black bodies and are not black at all. In usual experiments, a black body has a little hole that emits radiations which are analyzed by a spectrograph. The challenge which

took four decades to be resolved was to evaluate the spectrum of the cavity's radiation.

It was first discovered that the radiation spectrum didn't depend on the body radiation but only on its temperature T . It also proved that thermal radiation was an electromagnetic one. Hot objects like light bulbs and heated metals are close to black bodies. As the temperature increases, the black body color, corresponding to the spectrum peak shown in Figure 90, shifts from red to blue. There were various attempts to explain the black body radiation with thermodynamics and oscillators and to predict the spectrum curve.




Figure 90: Blackbody spectrum explanations over time. The sr in the radiance units decomposition is a steradian or square radiant, a unit of solid angle. T: temperature in Kelvin. k_B : Boltzmann constant. σ : Stefan-Boltzmann constant. h: Planck constant. c: speed of light and ϵ : emissivity. Compilation © Olivier Ezratty, 2021-September 2026.

Before Planck's work, Stefan-Boltzmann's law (1884) described the relation between temperature and total energy radiated per surface area ($\epsilon \sigma T^4$) and Wien's displacement law (1893) described the relationship between peak wavelength and temperature. These two laws worked well. **Wilhelm Wien** (1864-1926, German) was awarded the 1911 Nobel prize in Physics for this discovery.

Predicting the spectrum curve didn't work so well. First, Wien devised another law in 1896, Wien's approximation or radiation law that didn't work well with large wavelengths. The Rayleigh-Jeans formula created in 1900 didn't work for small wavelengths, leading to the so-called ultra-violet catastrophe. It was based on Boltzmann's statistical methods.

To make a better curve prediction, Max Planck guessed that the energy of the oscillators in the cavity was quantized and was a multiple of some quantity with the formula $E = nh\nu$, n being an integer, h being Planck's constant and ν the wave frequency. With this discretization, oscillators couldn't afford to have many energy quanta for high energy levels. Thus, their number decreased as the frequency increased instead of growing exponentially as in Rayleigh-Jeans law.

There was however no clear explanation of the origin of these quanta. The second step was Albert Einstein's work on the photoelectric effect in 1905, explaining how light and electrons interacted in quantized form. He guessed that the energy from an electromagnetic field is not spread over a spherical wavefront but is localized in individual directional quanta, which were later described as wave packets with a speed (of light) and length.

3.2.2 Electrons quantum numbers

The Niels Bohr's atomic model elaborated in 1913 helped describe the electron energy transitions within atoms and between orbitals that explained the various hydrogen emission spectrums experimentally discovered by **Johann Balmer** in 1885, **Theodore Lyman** in 1906 and **Friedrich Paschen** in 1908, corresponding to transitions ending respectively on the second, first and third atom electron layers. These are known as Balmer series, Lyman series and Paschen series. Other series ending on the fourth, fifth, sixth and seventh layers were later discovered in the infrared by **Frederick Sumner Brackett** in 1922, **August Herman Pfund** in 1924, **Curtis Judson Humphreys** in 1953, **Peter Hansen** and **John Strong** in 1973. The Bohr model was confirmed in 1914 with the **James Franck** and **Gustav Hertz** experiments, with mercury atoms in a vacuum tube being bombarded with electrons and with studying their energy loss and bounce effects.

Other energy transitions like those from the **Zeeman** effect could only be explained by the existence of other quantum numbers. During the 1920s, a better understanding of the quantum nature of electrons was achieved. It was progressively discovered that electrons in atom shells had four quantum numbers, as shown in Figure 91. **Friedrich Hund's** rules were created in 1925 to determine how electrons fill atom orbitals. It deals with the arrangement of electrons in degenerate orbitals, which are orbitals within the same subshell that have the same energy, unless some magnetic field is applied on the atom, generating the Zeeman effect and lifting this degeneracy. As the atoms grow in number of electrons, the valence shell orbitals are progressively

filled first with electrons with the same spins and when they are saturated, electrons with opposite spins fill the remainder of these orbitals. When a shell is filled, new electrons are added in the next shell, and so on.

The four electron quantum numbers are:

- **n principal quantum number** corresponding to their energy level or electron shell in the atom electron shells, numbered from 1=K, 2=L, 3=M to n , n being very high for so-called Rydberg (high-energy) states close to atom ionization. The principal quantum number is limited to 7 for non-excited known atoms and is theoretically illimited with excited atoms. A record of $n=766$ was observed with hydrogen atoms in interstellar medium. This number may correspond to some energy levels used in cold atoms and trapped ions qubits. It corresponds to the rows shown in Figure 92.
- **l orbital angular momentum** numbered from 0 to $n-1$ or letters (s, p, d, f, g, h, i, etc.) also named azimuthal or orbital quantum number, which describes the electron subshell and quantifies its amplitude. It corresponds to different types of orbitals around the atom and to the columns as shown in Figure 92.
- **m_l magnetic quantum number** labels the quantized projection of orbital angular momentum on a chosen axis. Its value or angle is also quantized and is an integer between $-l$ and l .
- **m_s spin projection quantum number** being either $+1/2$ or $-1/2$, in a given spatial direction (usually x, y or z in an orthonormal basis), called spin component, also named intrinsic angular momentum. This is the property used in so-called electron spin or silicon qubits. But...what is the unit of the spin? It is rarely mentioned but the spin unit is the Dirac constant $\hbar$, which equals the Planck constant h divided by 2π . It is an intrinsic property which doesn't depend on the situation like temperature. What physical property is it describing? It seems it doesn't describe a physical rotation of the electron around an axis. It may still be linked to some field rotation [349, 350].

The spin itself is represented by angular-momentum operators S_x, S_y, S_z . These different components generally do not commute and cannot all have simultaneous definite values. S^2 is the square modulus of the vector. S^2 and S_i (a projection of the spin on axis i) commute and are therefore measurable simultaneously. A projection on two axis S_i and S_j do not commute. A spinor is a vector representing the spin

of a $1/2$ spin particle using the notation $\chi = \begin{pmatrix} \chi_1 \\ \chi_2 \end{pmatrix}$,

with χ_1 and χ_2 being the complex components of the spin up and down along a given Z direction. Logically, it is like a spin qubit and represents the spin vector in a Hilbert space of dimension 2. Spinor trans-

forms belong to the special unitary group $SU(2)$, aka the Lie group of 2×2 unitary matrices.

It is key to understand the effect of these various electron quantum numbers in many fields like with NV centers and silicon spin qubits, quantum dot photon sources and many others.

A quantum system **ground state** corresponds to its lowest energy state among its possible discrete energy levels. Excited states correspond to higher energy levels that a quantum system can occupy when it absorbs outside energy. The ground state is described by a wave function which is usually less complex than those of excited states and often has unique symmetry properties.

Zero-point energy is a non-zero minimum energy of quantum systems in their ground state. This is a consequence of the Heisenberg uncertainty principle, which prevents the system from having precisely zero kinetic energy. The ground state is usually the most stable state of a quantum system. Systems naturally tend to lose energy and settle into the ground state unless they are provided with sufficient energy to move to an excited state. One way to prepare a system in its ground state is to use very low temperatures near 0 K. In quantum field theory, the ground state is referred to as the **vacuum state** which represents the state with the lowest possible energy, with no particles present.

Orbital types. Atom electrons can be in three types of orbitals (core, valence, excited, Rydberg which are highly excited states beyond $n \geq 30$), all being bound states, and continuum states (atom becomes ionized, electrons become free electrons, and their energy is not quantized anymore) (Figure 93). Inert elements have a closed shell configuration.

In Rydberg atoms, the core electrons shield the outer electron from the electric field of the nucleus such that, from a distance, the electric potential looks identical to that of the electron in a hydrogen atom. The radius r of the electron orbital scales as n^2 times the Bohr radius a_0 , which is the $n = 1$ value, n being the electron principal number, $k = 1/(4\pi\epsilon_0)$ aka k_e or the Coulomb constant, ϵ_0 being the vacuum permittivity, e is the elementary charge, and m is the electron mass.

The Rydberg states energy levels are spaced by $1/n^3$ and converge towards the ionization level.

$$r = n^2 a_0 \text{ with } a_0 = \frac{\hbar^2}{k e^2 m}$$

The orbital cross-section, whether circular or elliptical, scales as n^4 for hydrogen and can become quite large for high n . At $n = 100$, the radius $n^2 a_0$ already reaches about $0.5 \mu\text{m}$, so a cross-section of order $1 \mu\text{m}^2$. The Coulomb interaction is the force F between two charged particles. It is inversely proportional to the square of their distance and proportional to the product of their charge, which can be negative or positive.




Figure 91: The four electron quantum numbers explained visually. They are not intrinsic to the electron but rather to the behavior of electrons in atoms. Green labels correspond to experiments and black labels to theory © Olivier Ezratty, 2023-2026.




Figure 92: Visual representations of the electron atomic orbitals of the hydrogen atom, corresponding to their principal number n , orbital angular momentum ℓ and magnetic quantum number m_ℓ . Source: Keith Enevoldsen adapted by Olivier Ezratty, 2023.

$$|F| = k_e \frac{|q_1||q_2|}{r^2}$$

We've seen that the electron quantum numbers are the consequence of its interaction within the atom. These are driven by multiple physical laws:




Figure 93: A simple representation of the various states of an electron in an atom shell. They can be located in inner orbital shells of the atom, then in the valence shell which is the last of the shells for a non excited atom, then in other shells for excited states, and Rydberg states which create either elliptical or circular shapes. © Olivier Ezratty, August 2024.

- The **Coulomb force** which is the electrostatic attraction between the positive nucleus and negative electrons.
- The **Pauli exclusion principle** according to which no two electrons in an atom can occupy the same quantum state simultaneously. It derives from electrons being fermions, individual particles with a half-integer spin. As a result, electrons are forced into higher energy shells ($n = 1, 2, 3, \dots$) rather than sitting in the lowest energy level. The consequence is the structure of the periodic table of elements.
- The **Heisenberg indeterminacy principle**. The more precisely you know a particle's position (Δx), the less precisely you can know its momentum (Δp), such that $\Delta x \Delta p \geq \frac{\hbar}{2}$. If an electron was falling into the atom nucleus, its position would be extremely localized (Δx becomes very small). This would cause its momentum (Δp) to be very large, giving it enough kinetic energy to be expelled from the atom. The Heisenberg indeterminacy principle is derived from the wave-particle duality nature of the electron. It is a mathematical property of waves. If you represent a particle as a wave, its position and its wavelength (related to momentum) are related by a mathematical operation called a Fourier Transform. It is mathematically impossible for a wave to be both a sharp spike (precise position) and a single frequency (precise momentum) at the same time. The narrower the spike, the more frequencies you need to build it.
- The **spin-orbit interaction** which is a relativistic effect coupling the intrinsic electron property (spin) with the orbital motion properties. From the electron's perspective, the nucleus is orbiting it. This moving positive charge creates a magnetic field. The

electron's magnetic moment coming from its spin interacts with this magnetic field. This shifts the energy levels slightly, splitting what looks like a single spectral line into several fine lines (fine structure) in the atomic spectrum. The effect is described by the Dirac equation that combines Schrödinger's quantum equation with Einstein's special relativity.

- **Quantum electrodynamics (QED) and vacuum fluctuations**. In the empty space between the nucleus and the electron, interactions occur with the quantized electromagnetic field which shows up as virtual photons popping in and out continuously of the vacuum. The electron's interaction with its own electric field through these fluctuations produces the Lamb shift, a tiny but measurable displacement of certain energy levels.

By the way, let's mention that the mass of the electron can be precisely measured using Penning traps. It works with trapping a single electron in a magnetic and electric field and measuring how fast it travels continuously in a circle in that field, *aka* "cyclotron frequency". The cyclotron frequency is $\omega_c = eB/m$, so the electron mass is $m = eB/\omega_c$, e being the elementary charge, B the magnetic field and ω_c the cyclotron frequency, which is measured thanks to the oscillating voltage detected on the Penning trap metal plates [351].

3.2.3 Nucleons and nucleus quantum properties

Electrons have a spin but also atom nucleus and their nucleons constituents, which are neutrons and protons, themselves made of quarks. An atom nucleus has Z protons corresponding to the element atomic number and N neutrons which add up to a total of $A = Z + N$

nucleons. Protons and neutrons have a spin of $1/2$. The nucleus has a half-integer spin ($1/2, 3/2, 5/2, \dots$) when the number of neutrons plus the number of protons is odd, an integer spin ($1, 2, 3, \dots$) when the number of neutrons and protons are both odd and no spin at all when its number of neutrons and protons are both even.

Nucleus spins are either something we need to avoid like in silicon qubits produced with ^{28}Si , the silicon isotope with a null spin, or that we use to store qubit information like in NV centers and SiC cavities and electron donor qubits based on atoms like phosphorus where there is a coupling between some atom nucleus spins and some free electrons.

How is it possible to have a zero spin when you add up the spins of protons and neutrons which are positive? Let's take a pause and provide some answers. This is due to the way a nuclear spin is calculated.

Nucleons are composite quantum objects made of quarks that have quantum properties similar to electrons quantum numbers, but with different possible value bounds and meanings:

- **n nucleon shell number** or layer with integer values ranging from 0 to 6. It is bounded and there is no equivalent of Rydberg states in nucleus with large principal quantum number. The nuclear shell model is the equivalent of the atomic Bohr model related to electron shells.
- **ℓ orbital angular momentum** which is also quantized with an integer value starting at 0, the angular momentum itself being $L^* = \hbar\sqrt{\ell(\ell+1)}$. ℓ is labelled s, p, d, f, g, h, i like with electrons.
- **m_ℓ magnetic quantum number** with integer values ranging from $-\ell$ to ℓ . In each nucleon shell, nucleons of the same type tend to regroup by pairs with opposite magnetic quantum numbers.
- **m_s spin quantum number**, being either $+1/2$ or $-1/2$, in each spatial direction, also named intrinsic angular momentum. A nucleon spin s equal to $1/2$ is the size of the vector $\vec{s}$. The spin angular momentum is $S^* = \hbar\sqrt{s(s+1)}$ with $s=1/2$.

A nucleon total angular momentum is a vector $\vec{j} = \vec{\ell} + \vec{s}$ and $j = \ell + s$ in scalar representation with ℓ being the nucleon orbital angular momentum and $s = 1/2$ its intrinsic angular momentum or spin. The scalar representation is a good approximation of the vector representation since nucleons move in an average magnetic field orienting them in a similar direction. In the end, the atom **nuclear spin** is the sum of its nucleon's total angular momentum j .

As atomic nucleus size grows, nucleus shells are filled progressively. The first atom nucleus shell model was proposed by **Dmitri Ivanenko** with **E. Gapon** in Russia in 1932. The model was developed in 1949 following

independent work by several physicists, including **Eugene Paul Wigner**, **Maria Goeppert Mayer** and **J. Hans D. Jensen**, who shared the 1963 Nobel Prize in Physics.

Filled layers have a number of neutrons or protons called “magic numbers” (2, 8, 20, 28, 50, 82, 126) as shown in Figure 94. Atoms with entirely filled layers of either neutrons or protons are more stable. Nucleons pair in nuclei orbits with projections $\pm m_\ell$ such that their momentum cancel. The notion of layer magic number applies separately for protons and for neutrons. For example, ^{116}Sn (tin) has a magic number of 50 protons and ^{54}Fe (iron) has a magic number of 28 neutrons. Filled shells have a total angular momentum of zero since made of pairs of neutrons or protons with opposite projections of total angular momentum. That's why an even number of protons and neutrons lead us to have a zero nuclear spin. And when you have both a magic number of neutrons and protons, your nucleus is doubly magic like with ^{40}Ca and ^{208}Pb and has exceptional stability. All this refers to atoms and nucleus in their ground state.

On top of these numbers, nucleus have a parity that is $\lambda = (-1)^\ell$ where ℓ is the total orbital angular momentum of the nucleus. Its value corresponds to the symmetrical or asymmetrical structure of the nucleus.

The cohesion of the atom nucleus comes from the strength of the nuclear force that binds nucleons together. It is countered by Coulomb's force that creates a repulsion between same charge particles like protons. The relative value of the nuclear force and the Coulomb repulsion force explain nuclear fusion for small atoms and fission for large atoms. Iron sits at the maximum of binding energy per nucleon, around 8.8 MeV, which is why neither fusion nor fission releases energy beyond it, and which drives the iron abundance peak. This iron peak is explained by the fact that lighter elements were created by ordinary stellar nucleosynthesis and heavier elements by explosive nucleosynthesis in star supernovas.

Other notions are related to atom's nucleus and are frequently mentioned elsewhere in this book:

Hyperfine structure relate to the small differences in otherwise degenerate (equivalent, equal) energy levels in atoms, molecules and ions that are explained by the electromagnetic multipole interaction between the nucleus and electron clouds. In atoms, hyperfine structure comes from the energy of the nuclear magnetic dipole moment interacting with the magnetic field generated by the atom electrons and the energy of the nuclear electric quadrupole moment in the electric field gradient due to the distribution of charge within the atom.

Angular momenta coupling is the combination of the angular momenta of different systems or particles, resulting in a new total angular momentum that has significant implications for the physical properties and

Scattering Experiments provides information about the nuclear spin and magnetic properties of atoms by analyzing how the scattered particles interact with the atomic nuclei.

3.2.4 Photon quantum numbers

A photon is one quantum of excitation of a mode of the electromagnetic field. In free space, that mode is a plane wave whose electric and magnetic field variations are orthogonal to each other and to the propagation direction. The transversality is a property of the mode, not of some little object: in a waveguide, a cavity or an evanescent field, the mode shape and the field orientations differ.

Photons are described by their own quantum numbers which are:

- **Mass and electric charge** equal to zero. However, photons may have a tiny residual mass according to some theories [355]. There is even an upper bound for this potential photon mass, at $2 \times 10^{-24} m_e$, m_e being the electron mass, corresponding to 10^{-18} eV. It would represent a tiny portion of a photon energy, which is in the eV range for optical photons.
- **Wavelength** λ or frequency ν determines the photon energy $E = h\nu$ and momentum $p = h/\lambda$. Two distinct concepts associated with these properties can easily be confused. A Fock state has a definite photon occupation number in a specified spatial, spectral-temporal, and polarization mode. A single-photon wave packet is a one-photon Fock state whose mode has a finite spectral bandwidth, described by an amplitude $f(\omega)$ with well-defined relative phases between its frequency components. The photon number specifies how many photons occupy the same mode, while the wave-packet envelope specifies the spectral and temporal shape of that mode. Femtosecond lasers generate broadband optical pulses, mostly in the visible and infrared ranges, although these are usually coherent states rather than single-photon Fock states. Similarly, digital-to-analog microwave generators produce shaped coherent microwave pulses, such as those used to drive superconducting and electron-spin qubits. Shaped single-photon wave packets can instead be emitted by appropriately controlled quantum emitters, cavities, or qubits.
- **Spin angular momentum (SAM)**. The photon spin quantum number is $s = 1$, which makes it a boson. What is measured is its projection on the propagation axis, called **helicity**, and it takes only two values, $+\hbar$ and $-\hbar$ ($m_s = +1$ and -1 , spins being expressed in $\hbar$ units), corresponding to right and left circular polarization. A massive spin 1 particle would have three projections, including $m_s = 0$. The photon is missing that third one precisely because it is massless and travels at c : there is no rest frame in which the spin could be reoriented, so the longitudinal component has no physical existence. This is why photon

polarization lives in a two-dimensional Hilbert space and behaves like a qubit even though the photon has spin 1. The two circular states form a basis for that space, so *any* polarization state, including horizontal and vertical linear polarization, can be written as a linear combination of them, and conversely a circularly polarized photon is a superposition of horizontal and vertical.

- **Orbital angular momentum (OAM)** where the electromagnetic field wavefront is rotating helically along its propagation axis or vector [356]. It equals $\ell\hbar$ with ℓ being an unbounded integer, and it is a genuinely quantized property, not a continuous one. It corresponds to the photon spatial distribution perpendicular to its propagation direction. Because ℓ is unbounded, OAM spans a Hilbert space of arbitrarily large dimension for a single photon, which makes it a natural carrier for qudits rather than qubits.

Photons interact with other particles, mainly electrons either tied to atom nucleus, for photons absorption and/or emission, or free electrons like with the Compton effect.

They can be created, destroyed and modified by many of these interactions. Pairs of photons can also be generated by the annihilation of particles with their antiparticles. The classical field and the shape of the modes are described by Maxwell's equations, but the creation and destruction of photons and everything that counts them one by one requires quantum electrodynamics, the quantized version of that same field theory.

Photon directionality. Is a photon directional? Textbooks usually make a distinction between spontaneous emission with photons going in any direction, from a lightbulb or the Sun, and stimulated emission, with directional light, coming from lasers. Radio-frequency antennas can also create spherical radiation going in many directions.

But whatever its source and wavelength, a single photon is mostly always directional and moving in space as a planar wave. A photon electromagnetic wave is represented by orthogonal electric and magnetic fields variations travelling along a vector orthogonal to them. A photon direction can change when it traverses various materials having different refraction indices.

Can we have non-planar photons? “Any direction” photons can come from a statistical view of random multidirectional photons emissions or from a coherent superposition of photons emitted in several directions. With light bulbs, many photons are emitted in various directions by random thermal processes, with various photon wavelengths. Laser coherent light is made of photons with the same wavelength, phase and direction.

The distinction between a wave and a point-like particle is as blurred with photons as it is with electrons as far as their exact physical nature and dimensional scope is concerned (Figure 97).

Photon length and size are notions that are rarely mentioned, and for a good reason: in the standard quantum electrodynamics picture, a photon has no intrinsic size. It is a quantum of excitation of a field mode, and it is the *mode* that has a spatial extent, not the photon. The same photon energy can therefore come in a packet a few micrometers long or a few kilometers long, depending only on how it was made.

What defines that extent is the spectral width of the emission process. A wave packet of bandwidth $\Delta\nu$ has a coherence time $\tau_c \approx 1/\Delta\nu$ and a longitudinal length $L \approx c\tau_c = c/\Delta\nu$. Three orders of magnitude give the range. An atomic dipole transition with a 1 ns excited-state lifetime has a natural linewidth near 160 MHz, so its photon is a packet about 30 cm long, which is 500,000 wavelengths at 600 nm. A Quandel quantum dot single photon of 150 ps duration is about 4.5 cm long. A 10 fs pulse from a titanium-sapphire laser is 3 μm long, only a handful of wavelengths. None of these lengths is the wavelength, and the ratio between the shortest and the longest is around 10^5 .

A few non-mainstream proposals do assign a fixed geometric size to the photon. The Hunter-Wadlinger soliton model published between 1985 and 1989 describes an optical photon as an elongated ellipsoid of length λ and diameter λ/π [357], and other authors propose a longitudinal length of $\lambda/2$ [358, 359]. These circulate mostly outside mainstream peer-reviewed optics journals and are not part of the accepted quantum optics framework. They are mentioned here because readers do run into them online, not because they are established.




Figure 95: Photon wave packet or pulse. It has many harmonic frequencies shaped like a Gaussian curve in their Fourier transform.

What this means is that the usual graphic representation in Figure 96 in green is not realistic. It shows an infinite monochromatic wave, which no physical emission process can produce. Every real photon is a packet with a beginning and an end, and therefore a spread of frequencies.




Figure 96: Single frequency photon... of undetermined length. Its photon Fourier transform is a single point. It is physically impossible to create such monochromatic photons. Since photons are emitted by physical emission processes like an electron dropping to a lower energy state in an atom, which happens over a finite duration, the resulting photon is always a finite wave packet.

The wavelength range is itself very broad, from below 1 nm for X-rays and several orders of magnitude smaller for gamma rays, up to over one millimeter and several kilometers for radio waves. It is different from the packet length just discussed: at 1550 nm, a telecom photon filtered to a 1 nm bandwidth has a packet length of about 2.4 mm, or some 1,500 wavelengths.

Whatever the packet length, its Fourier transform cannot be a single frequency. The conjugate pair at play here is not “photon length and wavelength” but the longitudinal position and the wavevector, $\Delta z \cdot \Delta k \geq 1/2$, or equivalently duration and energy, $\Delta t \cdot \Delta E \geq \hbar/2$. The shorter the packet in space or in time, the broader its spectrum must be, which is exactly the trade-off quantified in the time-bandwidth table further down [360]. A perfectly monochromatic photon would need an infinite packet, and therefore an emitter that has been running forever.

Photon modes correspond to properties of photons, with several types: spatial modes (where is the photon in space), temporal modes (when was the photon emitted), spectral modes (frequency or wavelengths) and orbital angular momentum (photons with a helical wave front) [361]. We find multimode photons in various quantum optics setups like with boson sampling experiment that we will describe a bit later, as well as in quantum key distribution settings [362]. Modes are formally defined by coherence and orthogonality properties of the EM field. These are orthogonal solutions of the EM wave equations. Different photon modes do not interfere. The energy of a linear superposition of modes equals the sum of the energy of the individual modes. Only photons with the same mode can be coherent and interfere.

Photon modes are the normalized and mutually orthogonal solutions of Maxwell’s equations for given boundary conditions (free space, waveguide, cavity), photons being the quantized excitations of these modes [361, 363]. They cover several degrees of freedom: spatial modes (the photon path or its transverse profile), polarization modes, and time-frequency modes (the photon temporal envelope and its spectrum, which are linked by a Fourier transform). Orbital angular momentum (OAM) modes, with their helical wavefront, are a particular family of transverse spatial modes [364]. The mode decomposition depends on the chosen basis. A single-photon wavepacket is single-mode in a basis built around its own shape, but multimode in a basis of monochromatic plane waves. Multimode quantum states of light are found in various quantum optics setups like the boson sampling experiments that we will describe later, as well as in quantum key distribution (QKD) [362]. The energy of a superposition of orthogonal modes equals the sum of the energies of the individual modes. Orthogonal modes do not interfere in the total detected power. They can still interfere when a linear optical element like a beam splitter or a polarizer projects them onto a common output mode. This is how a single photon in a superposition of two paths interferes with itself in a




Figure 97: Photon characteristics and polarization.

Mach-Zehnder interferometer (MZI). With two photons, as in the Hong-Ou-Mandel (HOM) effect, the interference visibility equals the squared overlap $|\langle\psi_1|\psi_2\rangle|^2$ of their modes in all degrees of freedom other than the one being mixed, and reaches 100% for indistinguishable photons [365].

Photon number is a way to describe groups of similar photons. Several photons with the same wavelength and polarization can be at the same place and at the same time. They also share the same direction vector. These photons are indistinguishable. This is a property of bosons, which are elementary or composite particles that can share the same quantum state, following Bose-Einstein statistics, while fermions cannot share a quantum state, following Fermi-Dirac statistics.

A group of similar photons in the same mode carries an energy $E = \hbar\nu N = \hbar\omega N$ above the vacuum, i.e. the energy of each photon multiplied by the number of photons in that mode. Counting the vacuum contribution, the total is $\hbar\omega(N + 1/2)$. Here, ν is the photon frequency, h in Planck's constant, $\hbar$ is Planck's reduced constant or Dirac's constant and ω is the photon angular frequency with $\omega = 2\pi\nu$, in radians per second, since one full oscillation cycle covers 2π radians.

A photon number counts the photons "clustered" in one mode, forming a higher energy EM field than a single photon. A powerful radio or digital TV emitter, and likewise a radar, produces a very large number of photons in the radio wave range, but the state it generates is a *coherent* state with an indefinite photon number, not a Fock state. True Fock states with a well defined N are, on the contrary, among the least classical states of light, as their Wigner function shown later makes clear. You can even create superpositions of multi-photons (or single-mode Fock state as we'll see later) with 0 (aka "vacuum state"), 1, 2 and 3 photons.

This can be used to create photon-number Bell states, namely entangled states of superposed photon numbers [366–368], leading to the teleportation of Fock states [369]. This is mind blowing and quite hard to visualize!

Quantum optics is heavily based on the model of the **quantum harmonic oscillator**, the quantum-mechanical analog of the classical harmonic oscillator, with quantized energy levels.

The energy of a quantum oscillator can be described with a simple equation, N being the photon number. When $N=0$, the oscillator energy corresponds to the vacuum state energy. With energy E and photon number N : $E = \hbar\omega(N + \frac{1}{2})$. The related photon wavenumber is $k = \frac{2\pi}{\lambda}$.

OAM control and readout. A photon orbital angular momentum $\ell\hbar$ can be modified or generated with various means. Generation can be done with a spiral phase plate which is a transparent plate whose thickness grows with azimuthal angle and depends on the photon wavelength, a liquid-crystal spatial light modulator (SLM, also used to control the tweezer light that controls atoms in cold-atom quantum computers), or with a patterned liquid-crystal element *aka* q-plate [370–372]. OAM measurement can use projective filtering with a hologram and single-mode fiber, interferometric sorting using a Mach-Zehnder interferometer with a Dove prism, a log-polar mode sorter, with single-shot intensity diagnostics, and with wavefront twistors [373–375]. You can even entangle photons with their OAM property [376].

Photon wavenumber is the spatial frequency of a wave, measured in radians per unit distance. It is defined as k with the above right formula using the photon wavelength.

Creation and annihilation operators or ladder operators are mathematical operators used with quantum harmonic oscillators and many-particles systems. An annihilation operator $\hat{a}$ reduces the number of particles in a given state by one and a creation operator $\hat{a}^\dagger$ increases this number by one. It is the adjoint operator of the annihilation operator. These operators act on states of various types of particles, and with photons, adding or removing a quantum of energy to an oscillator system.

Quantum optics applies the same quantum postulates to quantized field modes, typically using Fock space and creation/annihilation operators. The single-particle position-space Schrödinger equation is nonrelativistic and works with massive particles. However, the Schrödinger picture and equation $i\hbar\partial_t|\psi\rangle = H|\psi\rangle$ also works with field-mode states.

Mathematically, a photon occupation number operator is a Hermitian operator $\hat{N} = \hat{a}^\dagger\hat{a}$. And a photon num-

ber of n superposed photons created by the operator $\hat{a}^\dagger$ applied n times to the vacuum state $|0\rangle$ creates the state $|n\rangle = \frac{1}{\sqrt{n!}}(\hat{a}^\dagger)^n|0\rangle$. It is also possible to create a superposition of photon-number, expressed as $\sum_n c_n|n\rangle$.

Second quantization is the broad field of quantum physics that deals with many-body quantum systems. It was introduced by **Paul Dirac** in 1927 and developed afterwards by **Vladimir Fock** and **Pascual Jordan**. While the first quantization dealt with individual quantum objects and their description by the Schrödinger wave equation, the second quantization describes many-body systems which are represented mathematically by Fock states and Fock spaces.

Its formalism introduces creation and annihilation operators to construct and handle the Fock states, providing the mathematical tools to study quantum many-body systems.

Instead of describing such a system as a tensor product of all its constituent quantum objects, it is simplified with chaining $|n_{k_i}\rangle$ describing the n_i quantum objects that are in the same quantum state k_i as described in the above equation related to the creation operator.

A many-body system is described as the tensor product of the Fock states $|n_{k_i}\rangle$ corresponding to each individual quantum states in the system: $|n_{k_0}\rangle \otimes \cdots \otimes |n_{k_n}\rangle$, given that the occupation number n_i for the Fock state $|n_{k_i}\rangle$ can be 0 or 1 for fermions and any non-negative integer for bosons. When all occupation numbers are equal to zero, the Fock state corresponds to the vacuum state.

A Fock state with only one non-zero occupation number is a single-mode Fock state. Contrarily, a multi-mode Fock state has several non-zero occupation numbers.

Quantum optics. This field of quantum physics started quite late. In 1956, the Hanbury-Brown-Twiss (HBT) experiment was about observing the intensity correlations of the radiation of a mercury lamp and from some bright stars. After traversing a beamsplitter (with a mercury lamp) or at two spatially separated points (for stars), the intensities measured by two detectors were fluctuating, and these fluctuations were correlated.

It was then explained by the emission of photon bunches coming from thermal sources. But it could be explained without using photons and quantum physics.

Quantum optics really started when it became possible to create non classical light sources like pairs of photons and single photons, respectively in 1967 and 1977. Photon pairs were first created using cascaded atom decay and parametric down conversion [379].

Semi-classical light. It describes interactions between quantized matter such as atoms and electrons with classical light fields. Continuous laser light belongs to this category.

Non-classical light. Light and photons are always quantum, just because it comes from quantized energy

exchanges with matter. Still, light is considered to show non-classical and quantum effects when the electromagnetic field is quantized and photons are handled individually. This happens in a couple situations: creation of entangled Bell states, antibunching, photon noise and negative probabilities with the Wigner function. We'll look at each of these phenomena.

Bell states where single photons behave probabilistically and in the general case have no *a priori* properties like polarization, wavelength, wavevector before being measured. These properties related to entanglement are revealed while being measured and show correlations between entangled photons whose measured properties will be random.

Antibunching corresponds to a light field where two photons are less likely to arrive together than in a coherent laser field. It is defined by the second-order correlation at zero delay, $g^{(2)}(0) < 1$, an ideal single photon source reaching $g^{(2)}(0) = 0$ since it never emits two photons at once. A coherent state gives exactly 1 and a thermal source gives 2. It was first observed in 1977 by H. Jeff Kimble, Mario Dagenais and Leonard Mandel on the resonance fluorescence of a sodium atomic beam [380]. It is detected with a HBT (Hanbury Brown & Twiss) intensity auto-correlator showing a dip instead of a peak at zero delay. Antibunching usually comes together with sub-Poissonian photon statistics, that is a photon number distribution narrower than a Poisson distribution (Figure 98), but the two are distinct properties: the first is about the timing of pairs of detections, the second about the counting statistics over a time window, and states exist where only one of the two holds.

It can be generated by single photon sources as well as from pulse mode lasers. A coherent state from a laser has Poissonian statistics generating random photon spacing and a thermal source light field has super-Poissonian statistics and yields bunched photon spacing. All these aspects belong to the field of statistical optics.

The quality of single photon sources is measured with data from two experiments.

The first experiment uses a variant of a **Hanbury Brown and Twiss** (HBT) intensity autocorrelator that checks the photons are emitted in a very regular way, like a metronome. From a starting click on one of the two photon detectors, it analyzes the time distribution of the appearance of the following photons. This produces the plots in Figure 99. The ideal model would be that of a high peak on either side of the center. The low peaks represent the system noise. This experiment, originally created to detect the size of stars, also enabled the validation of the corpuscular nature of photons. The experiment can be interpreted intuitively: photons hit a balanced (50/50) beam splitter and are transmitted or reflected at random. Behind this mirror are two photon counters, here with SPADs (avalanche diodes). The system identifies when photons are detected at the same

Q How can we represent photons?

Is a photon always moving in a straight direction? In vacuum and absent external forces, individual photons generally travel in straight lines. When a photon exits a laser, it travels in a highly collimated, straight-line direction that is generated by the process of stimulated emission. Each emitted photon has the same phase, frequency, and direction, resulting in a coherent beam with minimal divergence. But when photons encounter gravitational fields, like near stars or black holes, or materials with varying refractive indices, their paths can bend due to gravitational lensing or refraction.

Can a photon be spherical? When an atom emits a photon via spontaneous emission, the electromagnetic field associated with that photon can be described classically as a spherical wavefront, radiating outward in all directions from the atom. The emission is isotropic in the absence of external fields or constraints. But a single photon isn't a classical wave. It's a quantum excitation of the electromagnetic field, and its spatial description is more subtle. It doesn't literally spread out like a spherical shell, nor does it

split or hit multiple targets. Instead, its probability amplitude can be spread spherically, meaning there's a chance it could be detected in any direction, but once measured, it's found at a single point. The photon's probabilistic wavefunction can be spherical but in terms of physical shape or trajectory, a photon is not a tiny sphere, nor does it expand like a balloon. It's more like a ripple in the quantum field in a single direction, with a probabilistic orientation depending on its emitter.

What's the size of its perturbation in space and time? A photon is a quantum excitation of the electromagnetic field, often modelled as a wavepacket with a certain spatial and temporal spread. Its spatial extent can range from micrometers to kilometers depending on how sharply its frequency is defined. A narrow frequency wavepacket has a long wave-packet in space and time. In quantum optics, single photons generated in experiments often have temporal durations on the order of femtoseconds to nanoseconds, corresponding to spatial lengths of microns to meters [377, 378].




Figure 98: Poissonian, sub-Poissonian and super-Poissonian photon *number* distributions and the corresponding photon arrival time distributions. Compilation: Olivier Ezratty, 2021.

time by both sensors. If the photons use the same path to reach both detectors, there will be no coincidence since the emitted photons are sent in well-ordered trains and can only be on one side or the other. By adding a delay line between the mirror and one of the sensors that is proportional to the period of emission of the photons, it creates many occurrences with photons arriving

simultaneously in both sensors. This is what we see in the two curves, one of them being with a linear scale of coincidences (measured over a period of time sufficient to capture hundreds of them) and another logarithmic which helps better characterize the noise of the system.

The second experiment, called HOM (Hong–Ou–Mandel) after **Chung Ki Hong**, **Zhe-Yu Ou** and **Leonard**

Mandel and published in 1987 [365], validates that the emitted photons are indeed identical and impossible to distinguish [381]. It sends the two photons onto the two input ports of a single balanced beam splitter, not a Mach-Zehnder interferometer. When they are perfectly indistinguishable, the two amplitudes leading to one detector on each side cancel out and both photons always leave through the same port, which shows up as a dip down to zero in the coincidence rate as the relative delay is scanned. The residual depth of that “HOM dip” measures the indistinguishability.

Quadratures representation is a way to describe the electromagnetic field and its related uncertainty. An EM wave is positioned in two axis X and Y or X_1 and X_2 corresponding to the rotation of the electric field in the EM field, thus the equations describing X_1 and X_2 below, with the cosine and sine of ωt (angular frequency $\times$ time). Said otherwise, a quadrature describes the real and imaginary parts of a complex amplitude. This EM field complex amplitude is rotating so what’s interesting is not the grey circle position in the chart but its shape and size which represents the photon measurement uncertainty.

It is represented by the variation of the length of the vector which is the photon number and of the width of the circle, orthogonally to the vector, which corresponds to the phase uncertainty. For unsqueezed coherent light, this uncertainty is the same as the vacuum state.

Photon noise aka shot noise is found in the detection of light and corresponds to quantum fluctuations in the electromagnetic field. This noise or imprecision can be squeezed in one dimension.

Squeezed light corresponds, in a quadrature or phasor diagram representation, to states whose uncertainty on one quadrature X_1 or X_2 is smaller than for the vacuum state, the excess being pushed onto the orthogonal quadrature so that the Heisenberg product is preserved. Depending on which quadrature is squeezed relative to the mean field direction, this translates into reduced amplitude (photon number) noise or reduced phase noise. It can be generated by different means like a parametric down conversion [383]. Balanced homodyne detectors are used to detect squeezed light.

Wigner function is yet another representation of a quantum state, richer than the phase diagram above which is used to measure the level of quantumness of a light pulse. It is not far from a probability distribution of the electric field in the (Q, P) plane that can take negative values in some conditions, for so-called non-Gaussian states [384–386]. With coherent states, $W(Q, P)$ is a symmetric Gaussian function peaking at the average values of the sine and cosine components of the electric field with the peak width corresponding to the vacuum noise like in the quadrature representation.

For a pure state in the position representation, the Wigner function equation looks like this, ψ being the

quantum object wave function, Q and P the position and momentum and y, the variable used in the integral:

$$W(Q, P) = \frac{1}{\pi\hbar} \int_{-\infty}^{+\infty} \psi^*(Q + y) \psi(Q - y) e^{\frac{2iPy}{\hbar}} dy$$

It returns a real value that can be positive or negative. Q and P could be replaced by the sine and cosine components of the quantized electric field like in the phasor diagram (see also Figure 100).

In Figure 101 are presented a set of Wigner functions, ranging from the most classical to the most quantum fields. A is a coherent state (the vacuum state has a similar Wigner function, but centered around $P=0$ and $Q=0$), B, a squeezed state, C a single-photon state and D a Schrödinger’s-cat state. Figure 102 completes the landscape with 2D representations of these states and others. The projections or shadows of the Wigner function are the probability distributions of the quantum continuous variables Q or P. The Wigner function is a Gaussian function for A and B but takes negative values for the non-Gaussian strongly quantum states C and D [387]. These negative values vanish very quickly with decoherence.

Parametric down-conversion is a nonlinear optical process converting one photon of high energy into a pair of photons of lower energy. It is used to generate pairs of entangled photons.

Photons zoo. Figure 103 shows various photon states as a summary of this section. Random photons in spontaneous light coming from the Sun or a light bulb, and bright coherent beams from a laser, both admit a classical field description. Genuine “photon number” Fock states with a well defined N do not, and belong squarely to non-classical light.

Other forms of light described here are semi-classical or non-classical: photon number superposition, squeezed states where the precision is improved in photon number, amplitude or phase at the expense of the others, single-mode coherent laser light, wave packets created by pulse lasers or microwave coming from waveform generators used with superconducting and electron spin qubits, entangled photons used in QKD and photon qubits, and non-Gaussian states which are weird beasts too complicated to describe in a couple of words that are used to implement non-Clifford quantum gates with photon qubits when using continuous-variables photon qubits.

3.2.5 Elementary particles quantum numbers

We have inventoried relatively standard quantum numbers with electrons and photons and did show some commonalities between atoms, electrons and photons.

There are several other quantum numbers that characterize elementary particles and their composites within

antibunching measurement

$$g^2(0) = A_0 / A_{\text{side}} = 0.0028$$

- time correlation of second order or of intensity between pairs of photons with a zero delay.
- the closer to 0, the better.
- describes the level of noise in the single photon source.
- value for Quandela's photon sources in 2022.





Figure 99: How antibunching is measured. The metric is $g^2(0)$, the second-order coherence function or the second-order correlation function at zero time delay. It is an anti-bunching metric, ensuring that only one photon is emitted at a time by a single photon source. It is measured with dividing the A_0 area corresponding to zero coincidence, and the A_{side} area of the first intensity peak around the zero coincidence point. In 2022, Quandela's source had a $g^2(0)$ of 0.0028 according to tests done in Australia [382]. Sources: various, and Quandela for the bottom charts. © Olivier Ezratty, updated in September 2026.




Figure 100: What squeezed light looks like when using the quadrature representation. © Olivier Ezratty, 2023.

the Standard Model. These include isospin (found in quarks and hadrons like protons and neutrons), charm, strangeness, topness, and bottomness (which are flavor quantum numbers representing the net number of their

respective quarks), baryon number (defined as 1/3 of the difference between the number of quarks and anti-quarks) and lepton number (the difference between the




Figure 101: Various 3D representations of the Wigner function, for Gaussian (vacuum state, coherent state, squeezed state, thermal state) and non-Gaussian light which exhibit negative Wigner function values (Fock states, cat-states). Source: Figure 1.2 in [388]. Added in September 2026.




Figure 102: Representation of the Wigner function for some photonic states, vacuum state, squeezed vacuum state, Fock state with a number of identical superposed photons, cat-state, squeezed-cat state and GKP state. Non-Gaussian states can have negative Wigner function values. Source: [389].

number of leptons and antileptons). See the inventory of elementary particles quantum numbers in Table 3.

3.2.6 Spins everywhere

Spins are nearly everywhere in quantum physics, quantum technologies and with every type of qubit. You will discover many qubit modalities in detail later in the book. Let's make a quick recapitulation of how spins are involved in each of them. There are many spins involved like individual electron spins, composite objects spin for atoms or atom nuclei. An individual electron spin vector sits in a Hilbert space of dimension 2, like a photon polarization which is a spin 1 relativistic particle.

Cold-atom hyperfine structures arising from the interaction between the atoms nuclear spin and the electronic magnetic moment define the qubit states $|0\rangle$ and $|1\rangle$ and intermediate energy states. Coherent manipulation of

spin states is achieved using laser pulses, microwave or radiofrequency (RF) fields to drive transitions between different hyperfine levels.

Trapped ions also frequently use these hyperfine states or other states involving electron and nucleus spins of the trapped ions (Figure 104). Like with cold-atoms, spin-orbit coupling is also at play.

Silicon qubit manifolds are the spin orientation along a given axis defined by a magnetic field, of individual electrons trapped in potential wells. Qubit readout frequently uses a singlet/triplet dual-electrons configuration, with an ancilla electron and spin-to-charge conversion.

NV centers use the electron spin of the negatively charged NV^- defect, whose ground state is a spin triplet ($S = 1$). The qubit is encoded in two of its three m_s sublevels, usually $m_s = 0$ and $m_s = -1$, which are split by about 2.87 GHz at zero magnetic field. Metastable singlet states sitting between the ground and excited




Figure 103: A zoo of photons separating classical states of light from non-classical states. A wave packet is not non-classical by default. It is the case for single photon wavepackets. © Olivier Ezratty, 2021-September 2026.




Figure 104: Various energy transitions involving spins, in a trapped-ion setting. It uses two subspaces corresponding to different spins, 3/2 and 5/2, or energy levels. ps pulse means picosecond pulse. Source: [390].

triplets provide a spin-dependent decay path, which is what enables optical spin initialization and readout at room temperature. Singlet/triplet correspond to the energy states of two electrons in an atom like helium or trapped in a crystal lattice cavity. A singlet corresponds to the ground state and is a superposed and entangled state of opposite spin electrons. The triplets correspond to three other electron spin configurations, one with same-spin electrons but in different orbitals and a third one, with superposed states (Figure 105). NV centers also couple the electron spins in the crystal cavity with the nuclear spins of nearby non-zero-spin isotopes, ^{13}C ($I = 1/2$), ^{14}N ($I = 1$) and ^{15}N ($I = 1/2$).

Superconducting qubits are based on Cooper pair of opposite spin electrons, which are used to create superconducting anharmonic oscillators in cavities, *aka* artificial atoms. These artificial atoms are often presented as spins.

Photon qubits are frequently generated by single photon sources made with “III/V” atom ensembles where the properties of the generated photons are depending on the atom ensemble spin that is controlled by a magnetic field. In the Mendeleev table, group III and V atoms have respectively 3 and 5 valence electrons, al-

Quantum number	Applicable to	Description
Electric charge (Q)	All particles	Determines electromagnetic interactions, quantized in units of the elementary charge.
Spin (s)	All particles	Intrinsic angular momentum, distinguishes fermions (half-integer spin) and bosons (integer spin).
Parity (P)	All particles	Symmetry under spatial inversion, conserved in strong and electromagnetic interactions.
Isospin (I)	Quarks, hadrons	Analogous to spin, describes up/down quark symmetry in strong interactions.
Weak isospin (T_3)	Left-handed fermions, bosons	Governs weak interactions, part of electroweak unification.
Hypercharge (Y)	All particles	Related to weak isospin and electric charge via $Q = T_3 + \frac{Y}{2}$.
Strangeness (S)	Hadrons with strange quarks	Counts strange quarks, conserved in strong but not weak interactions.
Charm (C)	Hadrons with charm quarks	Counts charm quarks, with similar conservation rules as strangeness.
Bottomness (B')	Hadrons with bottom quarks	Also called beauty, counts bottom quarks.
Topness (T)	Hadrons with top quarks	Rarely used since top quarks decay before hadronizing.
Baryon number (B)	Quarks, baryons	+1 for baryons, -1 for antibaryons, conserved in most interactions.
Lepton number (L)	Leptons	+1 for leptons, -1 for antileptons, conserved in most interactions.
Flavor quantum numbers	Quarks and leptons	Includes strangeness, charm, bottomness, topness, and lepton family numbers.
Color charge	Quarks, gluons	Quantum number for strong interaction, comes in 3 types (red, green, blue).
Helicity	Massless particles	Projection of spin along direction of motion, relevant for neutrinos.
CP (Charge-Parity)	All particles	Combined symmetry, violated in some weak interactions (e.g., kaon decay).

Table 3: There are many other discrete quantum numbers in the standard model compared to the quantum numbers we inventoried with atoms, electrons and photons. We however won't mention them in the rest of this book. You may notice that this table doesn't include the orbital and magnetic quantum numbers from electrons and nucleons. They are not included in this standard model table related to elementary particles because they are specific to the behavior of electrons and nucleons in atomic physics, and not to elementary particle physics in general. Sources: Google Gemini, CoPilot, and some formatting.

lowing them to form a stable 8-electron “octet” when paired.

Spin-photon interfaces are also used in many situations, when the quantum state of an electron spin is mapped onto the state of a photon, particularly in the microwave regime. It is a technique used for spin qubit readout as well as in quantum networks [392, 393].

3.3 Wave-particle duality

We often read and hear that quantum objects like photons and electrons are both waves and particles. They behave differently depending on the way they are observed. In some experiments, these quantum objects

behave like classical waves, are not localized in space and generate interferences when added together, a bit like colors can mix (photons) and sounds can mix (acoustic waves). In other experiments, particle detections are localized events, and particle states carry energy and momentum.

It is usually impossible to observe these two behaviors simultaneously although there are some exceptions. Another interpretation is that quantum objects act as particles when observed and as waves when not observed and interact with each other. Various experiments such as Young's double-slit experiment show that both photons and electrons behave both as particles and as waves depending on the context and measurement system, generating interference fringes when observed as waves.




Figure 105: The different spin configurations for a singlet and triplet for two-spin configurations like in NV center cavities or in the helium atom. Source: [391].

You can observe the path of a quantum object or the interferences it creates, but not both simultaneously (Figure 106). This is the Bohr's principle of complementarity according to which it is not possible to apply observables simultaneously in terms of particles and waves.

It shows up in the Young slit experiment: if we let the quantum object traverse both slits, it behaves like a wave and creates interferences. If instead we determine which slit the quantum object went through, which is done by tagging the path with a marker such as an orthogonal polarization or an entangled ancilla rather than by closing a slit, the interference disappears. Closing one slit is a different and much less interesting experiment, since it removes the superposition altogether instead of revealing it and the quantum object behaves and is observed as a particle. The classical probabilities of particle observation do not add up to make for the observed interferences. This wave-particle duality is also interpreted thanks to the quantum physics mathematical formalism that relies on vectors that can add up linearly like waves. It led to the "which-way" problem. When interference fringes appear through the superposition of paths, quantum mechanics suggests that the particle does not take one path or the other, but rather exists in a state that samples both simultaneously until a measurement is made.

Wave-particle duality is used in many quantum computers to make physical qubits such as trapped ions interact with energy in the form of photons emitted by lasers. Qubits can also interfere with each other. It is also used in cold atom matter-waves in absolute gravimeters.


Historically, **Thomas Young**'s famous slit experiment did demonstrate that photons acted as waves. It was later formalized by **James Clerk Maxwell**'s electromagnetic equations in 1865. Photons particle behavior was first observed with the photoelectric effect discovered by **Heinrich Rudolf Hertz** in 1887 and later formalized by **Albert Einstein** in 1905 and confirmed by the Compton scattering experiment in 1923.

The electron wave behavior was theorized by **Louis de Broglie** in 1924 and shown experimentally in several steps, first, independently, by **George Paget Thomson** and by **Clinton Davisson** and **Lester Germer** in 1927 with a crystal diffraction setup, then by **Claus Jönsson** with a double slit experiment in 1961 and then, with experiments showing how single electrons could interfere with themselves, first by **Pier Giorgio Merli** et al. in 1974 and then by **Akira Tonomura** from Hitachi in 1989 [394]. In 2023, researchers in China could observe a wave-particle-superposition-like behavior of one photon or two photons [395]. Double-slit experiments with single photons interfering with themselves have also been implemented [396].

3.3.1 Schrödinger's wave equation

Wave-particle duality led Erwin Schrödinger to create his famous wave equation which describes massive non-relativistic quantum systems with a wave function defining probabilities of finding a particle at a particular position in space and time, as shown in Figure 107.

Here's how to understand the components of this equation and their implications:



photons acting as waves

experiment: **Thomas Young, 1801**


concept and equations: **James Clerk Maxwell, 1865**

photons acting as particles

experiment: **photoelectric effect, Heinrich Rudolf Hertz, 1887**

theory: **Albert Einstein, 1905**

experiment: **Compton scattering effect, 1923**



electrons acting as waves

concept and equation: **Louis de Broglie, 1924**

experiments:

G. P. Thomson, C. Davisson & L. Germer, 1927 (crystal diffraction),

Claus Jönsson, 1961 (double-slit experiment),

Pier Giorgio Merli et al, 1974 (single electrons, 'photos'),

Akira Tonomura, Hitachi, 1986 (single electrons, 'movie')

Figure 106: Wave-particle duality with photons and electrons theory and experiments. © compilation Olivier Ezratty, 2023.

i = imaginary number, its square equals -1

first derivative of the wave function vs time

second derivative of the wave function vs position (Laplacian)

potential energy function depends on the particle, its physical constraints and its position

$\hbar$ = Dirac constant

$\hbar = \frac{h}{2\pi}$

h = Planck constant

m = particle mass

total energy of particle

massive non relativistic particle wave function defining its evolution in time and space, returning a complex number. this is the equation unknown variable

kinetic energy (« impulsion » observable)

potential energy (« position » observable) equals zero for a free particle

$$i\hbar \frac{\partial \Psi(x,t)}{\partial t} = -\frac{\hbar^2}{2m} \frac{\partial^2 \Psi(x,t)}{\partial x^2} + V(x) \Psi(x,t)$$

$$\hat{H} = -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + V(x)$$

Hamiltonian: operator applicable to the particle wave function $\Psi(x, t)$ to evaluate its total energy, is a unitary operator.

(cc) Olivier Ezratty, 2021

Figure 107: The famous Schrodinger's wave equation explained in detail © Olivier Ezratty, 2021-2025.

- Its **unknown** is the wave function of the particle $\psi(x,t)$ that describes its probabilistic behavior in space and time. x indicates the position of the particle in space, with one, two or three dimensions depending on its constraints, and t is the time. This function returns a complex number that encodes the wave amplitude and phase at a given location.
- The full Schrodinger wave equation illustrates the principle of **energy conservation**. The item to the left of the equation describes the total energy of the particle at a given time and place. The elements on the right are the particle kinetic energy and its potential energy. As indicated in the quantum physics postulates part of this book, the Schrodinger's Hamiltonian, which is a time-dependent unitary matrix operator, is expressed differently with photons and with relativistic massive particles.
- The **squared modulus of the wave function** gives the probability density of finding the particle at loca-

tion x at time t . For an electron, which is the most commonly analyzed particle with this equation, it is an indication of the probability of finding it at a given distance from the nucleus of the atom. Logically, as a result, the sum of the probability of finding the electron at a given position relative to the nucleus. The probability of finding the particle somewhere is equal to 1. This is called a normalization constraint (shown in Figure 108). One of its consequences is Max Born's rule that we will see later. The modulus of a complex number is the size of its vector. If $z = a + ib$, the modulus $|z|$ of z is thus the square root of the sum of the squares of a and b , also shown in Figure 108.

- It is a **partial differential** equation, i.e., it connects its components via derivative functions, in this case of first degree (a slope on a curve) and second degree (an acceleration). The particle wave function appears three times in the equation: to the left of

the equation with a first derivative on the time of the wave function, to the right with a second derivative on its position and with a simple multiplication with the function $V(x)$.

- The **potential energy of the particle** is defined by the function $V(x)$ which depends only on the particle position in space and its physical constraints, in particular electromagnetic ones. When a particle is free and moves without constraints, this function returns zero. This function $V(x)$ is the main variable of Schrödinger's equation.
- Schrödinger's equation is **analytically solved** in a limited number of cases such as for the electron of a hydrogen atom, a free particle, a particle in a potential well or box or a quantum harmonic oscillator. In the most complex cases, the resolution of the equation requires non-analytical methods, raw calculation and simulation. It is one of the fields of application of quantum simulators and gate-based quantum computers to solve the Schrödinger equation in cases where analytical methods are not available. Any micro or macro-object has a Schrödinger wave function, all the way to the entire Universe. But the equation only makes practical sense for nanoscopic objects.

$$\mathbf{z} = a + ib \quad |\mathbf{z}| = \sqrt{a^2 + b^2}$$

complex number module
size of its vector in two dimensional space

↓

$$|\psi(x, t)|^2$$

wave function module square
probability to find the particle at position x at time t .
If it's time independant, the system is in a stationary state.

↓

$$\int_{-\infty}^{+\infty} |\psi(x, t)|^2 dx = 1$$

integral of the propability to find the particle
in any position equals 1

Figure 108: Constraints of the Schrodinger's equation
© Olivier Ezratty, 2021.

- The equation is **linear** over time. This means, among other things, that any combination of solutions of the equation becomes a new solution of the equation. This makes it possible to decompose a wave function into several elementary wave functions that are called the “eigenstates” of the quantum object. They correspond to the different energy levels of the particle that are discrete when the particle is constrained in space, like the electrons in an atom. One can indeed in this case derive the notion of quantification of the particle states from the Schrödinger equation. The linearity of this equation has a lot of consequences like superpositions, entanglement as well as the no-cloning theorem.

- The operator which acts on the right side and accumulates the second derivative, and the potential energy function is called a **Hamiltonian**, which describes the total energy of the system. We find this expression in the quantum annealing calculation with D-Wave and with quantum simulators.
- This equation is a **general postulate** that has been experimentally validated in a very large number of cases. Its interpretation has given rise to much debate, namely, is it a simple probabilistic model or does it describe reality? It is dealt with by the field of quantum physics foundations.
- The generic Schrödinger equation is said to be **time dependent**. This equation is presented in various ways depending on the needs and annotations. The second derivative of the wave function on the position of the particle is sometimes presented with the nabla sign squared (∇^2).
- A nabla operates a derivative on a scalar or vector function. The ∇^2 operates a second derivative, also called Laplacian. The most concise form of Schrödinger's equation is below, with the Hamiltonian operator on the left ($\hat{H}$) and the energy operator on the right ($\hat{E}$), both applying to the particle wave function $\psi(x, t)$.

$$\left[-\frac{\hbar^2}{2m} \nabla^2 + V \right] \Psi = i\hbar \frac{\partial}{\partial t} \Psi \quad \hat{H}\Psi = \hat{E}\Psi$$

- The **time-independent** form of Schrödinger's equation applies to particles in a stationary state. In this version of Schrödinger's equation, the energy operator E is a simple constant, a real number:

$$\left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) \right] \Psi(\mathbf{r}) = E\Psi(\mathbf{r})$$

- For a system of identical particles, the Schrödinger wave function is either **symmetric** or **antisymmetric** under particle exchange. For two particles with coordinates r_1 and r_2 , $\psi(r_1, r_2) = \psi(r_2, r_1)$ for a symmetric state and $\psi(r_1, r_2) = -\psi(r_2, r_1)$ for an antisymmetric state. The first case corresponds to identical bosons, which have integer spin and can occupy the same single-particle quantum state. The second corresponds to identical fermions, which have half-integer spin and obey Pauli's exclusion principle: two identical fermions cannot occupy the same single-particle quantum state. This boson-fermion distinction follows more fundamentally from the spin-statistics theorem, while Pauli's exclusion principle follows from the antisymmetry of fermionic wavefunctions.
- The $\psi(x, t)$ function must be a **continuous function** defined everywhere in space. It is complex-valued, so it is not bounded between 0 and 1: what is required

is that it be square-integrable, meaning that $\int |\psi|^2$ over all space converges, so it can be normalized to 1. Note that $|\psi|^2$ is a probability *density* per unit volume and can exceed 1 locally, only its integral is capped. It also has a single value, even in the case of superposition. In that case, the $\psi(x, t)$ is a linear superposition of two Psi functions and is itself a psi function. A quantum superposition is just another wave function.

For a system with several quantum objects, the wave function describes the quantum system state, or quantum state. According to the Copenhagen interpretation of quantum physics, the wave function from the Schrödinger equation contains the best description possible of a quantum system.

If electrons and photons can both behave as waves, they have not the same wavelengths. Indeed, a photon with an energy of 1 eV (electron-volt) has a wavelength λ of 1240 nm (in the infrared spectrum) while an electron with the same energy has a much shorter wavelength of 1.23 nm (in the X-ray spectrum). This short wavelength explains why we use electron microscopes to probe matter with a much better resolution than light-based microscopes.

Quantum trajectories describe the time evolution of quantum systems and their quantum state, which can be solved using the time-dependent Schrödinger's equation [397, 398]. They are related to the **quantum jump** phenomenon, which describes the continuous change happening for example during the measurement of a quantum object [399]. The first quantum jump effect was observed in 1986 in Austria with barium ion decay-ing [400].

The **Schrödinger picture** and the **Heisenberg picture** are two equivalent formulations in quantum mechanics. In the Schrödinger picture, the state vectors represented by their wave functions evolve over time according to the Schrödinger wave equation, while the operators representing observables remain time-independent. Inversely, in the Heisenberg picture, the state vectors are fixed, and the operators evolve over time according to the Heisenberg equation of motion. Despite their differing approaches to time dependence, both pictures produce the same physical predictions, with the Schrödinger picture often being more intuitive for time-dependent problems involving state vectors, and the Heisenberg picture being advantageous for focusing on observables and their commutation relations. You often see this distinction in quantum chemistry quantum algorithms.

Dirac-interaction picture. aka interaction representation, is an intermediate representation between the Schrödinger picture and the Heisenberg picture. It is named after Paul Dirac, who introduced it in 1932.

Relativistic particles obey to Dirac and Klein-Gordon wave equations while photons are described with Maxwell's equations combined with a formalism coming from the so-called second quantization which re-

groups superposed photons, use photon numbers, and creation/annihilation operators.

Nonlinear Schrödinger Equation (NLSE) is the variant of the Schrödinger wave equation which introduces a non-linear term with g being a coupling coefficient

$$i\hbar \frac{\partial \psi(\vec{x}, t)}{\partial t} = \left(-\frac{\hbar^2}{2m} \nabla^2 + g|\psi(\vec{x}, t)|^2 + V(\vec{x}, t) \right) \psi(\vec{x}, t)$$

The NLSE appears in classical physics with nonlinear optics (in fiber optics), fluid dynamics, plasma physics and electric transmission lines, and in quantum physics in Bose-Einstein condensates (BEC), superfluidity and superconductivity, and in polariton condensates, which are quantum fluids of light (see 3.10.6).

3.3.2 Delayed choice experiment

John Wheeler proposed various thought experiments between 1978 and 1984 to determine if light chose its path by sensing the experimental devices. The Wheeler's delayed choice or which-way experiment asked the question: when does a quantum object decide to travel as a wave or as a particle?

It led to various experiments like the 1999 quantum eraser, but the most decisive experiment was conducted by a team of French researchers led by Philippe Grangier in 2006 as shown in Figure 109 [125]. They generated pulses of single photons with an NV centers source created by Jean-François Roch, a pioneer in this domain, that were sent through a first beam splitter (BS_{input}) and a delay line of 48 meters. Then, the two beams traversed a dynamic-controlled beam-splitter by electro-optical modulator driven (BS_{output}) by a quantum random number generator (QRNG). The experiment has been reproduced many times since then with many variations [401].

Finally, using two photon detectors (N_1 and N_2), the measurement generated either path-resolving statistics or interference statistics. It did not indicate that the photon possessed a definite wave-or-particle history. The QRNG clock was near the photon source, but the QRNG was positioned close to the dynamic beamsplitter. The experiment determined that the wave/particle behavior of the photons in the interferometer was dependent on the choice of the measured observable at the end of the photon journey, not the beginning. And even when that choice was made at a position and a time sufficiently separated from the entrance of the photons in the interferometer. Although it is still debated, it does not require a backward in time effect explanation [402].

Other more delayed-choice sophisticated experiments are regularly done. A Chinese team demonstrated a generalized multipath wave-particle duality implemented by a large-scale silicon-integrated multipath interferometers

Q Don't be confused between waves and waves...

Waves have different meanings in the context of quantum physics.

An electromagnetic wave is a physical phenomenon described by classical physics, specifically Maxwell's equations. It is made of oscillating electric and magnetic fields that propagate through space, carrying energy and momentum. These fields are real, measurable quantities, and the wave itself can be directly detected as visible light, radio waves, or X-rays. Its ontology is rooted in physical reality. The wave exists independently of observation and interacts with matter in observable ways.

In contrast, the Schrödinger wave function has a fundamentally different interpretation. It is not a physical wave in space but a mathematical construct that encodes the probability amplitude of a quantum system's state. The square of its magnitude gives the

probability density of finding a particle in a particular location in space and time. Unlike electromagnetic waves, the wave function is not directly observable. It represents knowledge or information about a system rather than a tangible entity. Its ontology is epistemic or statistical, depending on interpretation. Some physicists view it as a tool for predicting outcomes, others as a real but non-classical entity.

So, while both are called “waves”, the electromagnetic wave is a physical oscillation of fields, and the Schrödinger wave function is a probabilistic description of quantum behavior. Their similarities in form mask deep differences in meaning and physical significance. Finally, when quantum physics deals with wave-particle duality, the meaning of wave is the second one with its probabilistic description for massive particles described by Schrödinger's equation, and a wave in the electromagnetic sense with photons.




Figure 109: Delayed choice experiment and its quantum eraser. Source: [125].

[401]. A delayed-choice entanglement swapping experiment was also achieved in 2012, where entanglement is produced a posteriori, using two pairs of entangled photons, after the entangled particles have been measured and may no longer exist [403].

3.3.3 Large objects wave behavior

The wave-particle duality was verified with atoms in 1991 in interferometry experiments involving lasers and classical optics. A Young double-slit experiment was first carried out in Vienna in 1999 with fullerene molecules (C_{60} , formed of 60 carbon atoms as in Figure 110) by Markus Arndt, Anton Zeilinger and co-workers [404], later reviewed and extended to the C_{70} variant [405] and in 2012 with molecules containing 58 and 114 atoms,

the latter named $F_{24}PcH_2$ being made of fluorine, carbon, oxygen, hydrogen and nitrogen [406]. Figure 111 shows the shape of the molecule. In 2019, the same kind of experiment was done with a slightly more complex molecule, a polypeptide of 15 amino acids which serves as an antibiotic, gramicidin A1 [407].

In 2021, other experiments led to the creation of larger quantum objects, sized 100 and 140 nm, and cooled at ultra-low temperatures [408–410]. Nowadays, experiments are done with even larger systems made of thousands of atoms [411].

3.3.4 Photon's wave-particle duality

On the other hand, photons can behave under certain conditions like particles. When they reach an atom,




Figure 110: C_{60} fullerene molecule.





Figure 111: $F_{24}PcH_2$ made of fluorine, carbon, oxygen, hydrogen and nitrogen. Pc is a phthalocyanine macrocycle, a large aromatic ring system made of four isoindole units, themselves made of a benzene ring fused to a pyrrole ring (five-membered ring containing one nitrogen atom). Sources: [406] and [412].



they can transmit it some kinetic motion. This is what makes it possible to generate the somewhat counter-intuitive physical phenomenon of atoms laser cooling using lasers and the Doppler effect. Temperature is related to the movement of atoms in their gaseous, liquid or solid medium. Lowering the temperature means slowing down the movement of atoms.

A Doppler effect is used to do this. The moving atoms are illuminated with a laser whose frequency is tuned just below the energy absorption level of the atoms as explained in Figure 112. The atoms moving towards the light will absorb the photons because these have an apparent frequency that is higher than the absorption level. This reduces the kinetic energy of the atoms receiving the photon. The atoms moving in the other direction will not absorb them because the apparent frequency of the incident photon is then below the absorption level, so it is unable to change the energy state of the atoms.

Thanks to the random movement of the atoms in all directions, after a certain time, the overall temperature drops. This phenomenon slows down once the velocity of the atoms falls below a certain threshold, which explains the Doppler cooling limit, set by the balance between the cooling force and the random recoil of spontaneous re-emission, and equal to about $145 \mu K$ for rubidium 87.

These techniques are used to cool atoms to temperatures close to absolute zero. It is used to prepare cold atoms and trapped ions used in certain types of quantum computers, often in combination with magnetic and/or electronic traps to control the atoms position. Doppler-based measurement is also used to evaluate the speed at which stars and galaxies move away from each other and to evaluate the rate of expansion of the Universe. Other atoms laser-based cooling methods crafted to reach lower temperatures include Sisyphus cooling first proposed by Claude Cohen-Tannoudji in 1989 and using two counter-propagating lasers using orthogonal

polarization, evaporative cooling using magneto-optical traps (MOT) and optical molasses with 3D Doppler effect.

The record low temperature for ultracold polar molecules was reached in 2019 with 50 nK, achieved by researchers from JILA, the joint laboratory of NIST and the University of Colorado. It was obtained with laser cooling of a gas containing 25,000 potassium-rubidium molecules.

In an experiment performed in 1985 [413], Grangier, Roger and Aspect demonstrated how single photons can interfere with themselves. The single-photon character of the experiment was certified by anticorrelation measurements done with the same source and detection scheme. This was, however, not the overlap of classical waves building up a continuous bright-and-dark intensity pattern. Instead, the fringes emerged statistically from individual single-photon detections. Each photon was always recorded as a single click in one of the two output detectors. What interfered was not the photon itself as a classical wave but the probability amplitudes describing its two possible paths through the interferometer. The interference is carried by these amplitudes and becomes visible only in the detection statistics accumulated over many photons, not in any instantaneous field oscillation. So the experiment does not show a physical superposition of two distinct light waves but rather a modulation of the detection probabilities of a single quantum coherently split between the two arms. This is why we say the photon “interferes with itself”, the phrase referring to its probability amplitudes rather than to a classical field oscillation.

3.4 Superposition and entanglement

Superposition and entanglement are directly related to the wave nature of quantum objects and to the linearity of the underlying mathematical models expressed in quantum physics postulates.




Figure 112: Explanation of Doppler effect with photons, © Olivier Ezratty, 2021-2025.

3.4.1 Superposition

The strawman's version of superposition is that quantum objects can be simultaneously in several states or locations, such as the direction of electron spin, upward or downward, the linear polarization of photons, horizontal or vertical, an electron spin (Figure 113), or the frequency, phase or energy of an oscillating current in a superconducting loop crossing the barrier of a Josephson junction. It is not correct according to canonical interpretations of quantum mechanics. It is more related to quantum objects behaving as waves when not being measured.

Superposition is a mathematical consequence of quantum postulates and wave-particle duality. It results from the fact that a linear combination of solutions to the Schrödinger equation is also a solution to this equation.

Several quantum states of a given quantum object can be added together or superposed. Superposition explains the interferences obtained with electrons in the 1961 double-slit experiment.




Figure 113: Electron spin superposition, it is an idealized view given that the spin vector is rotating (Larmor precession) and the up and down vector represent a projection of the spin vector on a given direction. © Olivier Ezratty, 2021-2023.

A quantum object is not per se in a superposition of various states. It has a single and predictable quantum state described by a probability distribution of given observables in a given measurement context. Measuring this property can provide different values according

to the probability distribution. Also, a state is in a superposition relatively to a specified basis.

According to the Copenhagen interpretation of quantum physics, one should not try to give a physical meaning to superposition before any measurement. Superposition can happen with various weird situations, as we'll later see. For example, you can create superposition between several photon Fock-states, meaning, superposing 0 photon, 1 photon and 2 photons, or even photon frequencies. You can even superpose temperatures [414] and thermodynamic evolutions with opposite time arrows [415] which can challenge your willingness to visualize what's it all about!




Figure 114: "Diosi-Penrose collapse times" scale showing how quantum systems coherence times decrease exponentially with their size. Source: [416].

In quantum computing, superposition shows up with qubits, which can have an internal "value" corresponding to a linear combination of their basis states $|0\rangle$ and $|1\rangle$ with complex amplitudes instead of having one being one of the two possible values (0 or 1) as with classical bits. This mathematical view is expanded to N qubits

whose internal state is characterized by 2^N complex amplitude values. This contributes to the massive parallelism enabled by quantum computers. It looks like it should enable some exponential computing capacity but it is not the case. As we'll investigate later in this book, superposition alone is not sufficient to generate an exponential computing time speedup compared to classical computing equivalent solutions.

We also need entanglement and some specific quantum gates to really bring some exponential acceleration. In the case of a single quantum object, superposition is a combination of states corresponding to several exclusive states of an observable. Coherence refers to the phase relation between basis components. Environmental entanglement suppresses those relations in the subsystem's reduced density matrix. Decoherence is a phenomenon that destroys superposition, particularly with quantum measurement. Quantum systems coherence time depends on many parameters that are endogenous (the object inherent stability) and exogenous (noise coming from the environment). It seems as a rule of thumb that the larger the quantum object, the lower its coherence time will be, according to the so-called "Diosi-Penrose collapse times" scale [416] (Figure 114).

A different rule may exist with regards to the number of coherent and entangled objects in a quantum system [417] (Figure 115). We don't know yet how many quantum objects can be entangled. It is an overarching question that may be answered some day with large scale quantum computers supporting millions of physical qubits.

3.4.2 Entanglement

The simplest way to describe an entangled state of two quantum objects is to say these have a correlated state when being measured, whatever the distance between them. They form sort of a single object. You measure one quantum property of a quantum object, then with the other, and the related results will be correlated. This can be checked with tests done on a system repeatedly prepared in a similar way, using a so-called Bell test. But measurement correlations describe only one part of entanglement since this notion of a "single object" predates any measurement. Entanglement can be generated using the same property of quantum objects of the same type, of different properties for the same object type or even different properties of different quantum object types, as shown in Figure 116. Entanglement also shows up naturally between the spin of the electron and the proton of an hydrogen atom [418] or with the spins of the two electrons in an helium atom [391].

Entanglement can be described using a mathematical viewpoint based on superposition. The mathematical representation of a quantum system AB made of two subsystems A and B is usually the tensor product of the two subsystems, meaning, a large vector or matrix: $H_{AB} = H_A \otimes H_B$. We will describe the shape of the

matrices representing quantum systems a little later. In that case, the system AB can be described by or decomposed with its individual parts A and B . There are, however, situations where you can linearly combine several of these composite quantum states, which becomes a new quantum state.

Mathematical non separability. In many cases, such a composite state cannot anymore be mathematically decomposed as the tensor product of two states. The composite quantum state becomes mathematically inseparable. That is when entanglement shows up! Entanglement is a direct consequence of superposition applied to multi-object systems. Entangled quantum objects cannot be considered as separate objects. Measurements on entangled systems produce correlations that violate Bell inequalities, but no controllable signal or locally observable change is transmitted faster than light. This is the principle of the nonlocality of quantum properties that disturbed Einstein in 1935 and spurred his famous EPR paper with Podolsky and Rosen. Using qubit's representation that we'll describe later, classical entangled two-qubit states are Bell pairs, like $\frac{|00\rangle + |11\rangle}{\sqrt{2}}$ or $\frac{|01\rangle + |10\rangle}{\sqrt{2}}$. You see that they are a simple linear combination of separable states ($|00\rangle$ and $|11\rangle$ or $|01\rangle$ and $|10\rangle$). If you measure the first qubit in both cases, you have an even 50%/50% chance to get a $|0\rangle$ or a $|1\rangle$. When you measure the second qubit, you then have a theoretical 100% chance to get respectively the same value of the opposite values $|1\rangle$ or $|0\rangle$. But you can't decide in advance what is the first measurement outcome (on Alice's side). So, you observe some synchronicity between two measurements but there is no determinism with the first readout value. However, this situation is considered as being trivial. In Alain Aspect's 1982 Bell inequalities experiment, the measurement basis for the two entangled photons is both random and different, using polarizers with different angles. Specific angles are used to maximize non classical correlations with photons, 22.5° and 67.5° as shown in Figure 117.

σ and z-score. Bell inequalities violation experiments are checked statistically with a standard deviation of the measurement of S against theory σ_S and a z-score in numbers of σ that equals $(S - 2)/\sigma_S$, which flags the significance of the violation, the higher being the better. For example, 2σ corresponds to a $\approx 95\%$ and 3σ to a $\approx 99.7\%$ confidence. The further S sits above 2 and the smaller σ_S , the higher the z-score, so accumulating more coincidences directly buys significance. The Aspect January 1982 two-channel polarizers experiment generated a 46σ thanks to the accumulation of 50,000 coincidences [419], a record at the time.

No-signaling. Having two correlated random values without any hidden variable explaining it is described as the "no-signaling principle". Also, there is no statistical difference between a "first" or "second" measurement of entangled pairs, meaning Bob doing the measurement before or after Alice and Alice did not send any




Figure 115: Plot showing how far quantum macroscopicity could go with regards to the extensive size measuring the amount of quantum coherence in number of atoms (N_{ext}) and an entangled size quantifying the number of entangled particles (N_{ent}). The experiments covered are mechanical oscillators (filled triangles), levitated nanoparticle (filled diamond), trapped ion (filled square), diffracting molecules (filled circle) and hypothetical systems (unfilled triangles), including Leggett's crystal (unfilled squares). Source: [417].




Figure 116: Various scenarios for the creation of entanglement between different quantum objects of the same type or not, and between discrete and continuous variables corresponding to different degrees of freedom. One very special case is the entanglement between two degrees of liberty of the same quantum object, called hyper-entanglement. © Olivier Ezratty, 2026.

actual pre-determined information to Bob when doing the measurement on her side.

Creating entanglement. What we have described so far is a mathematical representation of entanglement. You might wonder how these composite objects are ar-

ranged in the real world. Some physical interaction must be created to entangle electrons, atoms and/or photons [420]. Photons can be prepared to be entangled with being generated by some excitement of atoms like calcium which generates photon couples of different

wavelengths but with some correlated properties like their polarization. Neutral atoms can be entangled with exciting them with lasers, raising their energy levels to a so-called Rydberg state, which then creates links with neighbor atoms that are within the radius of the excited Rydberg atom.

Electron spins are entangled with lowering a potential energy barrier between them. Quantum objects of different types can also be entangled, like photons with atoms or electron spins with photons. In 2017, researchers in Warsaw were able to entangle a photon with billions of rubidium atoms [421]. Also, the two electrons of a helium atom in the ground state are entangled.

These entangled quantum objects are not linked by chance. They usually had a common past or some past interactions. For example, two entangled photons can be produced by spontaneous parametric down-conversion in a birefringent nonlinear crystal such as BBO, then separated and routed by dichroic or polarizing beam splitters, giving two photons of orthogonal polarizations. The action on one of the two photons has an impact on the other photon as demonstrated by Alain Aspect in his famous 1982 experiment. But the values that are generated are completely random! It is not defined at one end but transmitted to the other end. It is a random value that can be uncovered at two different places with some quantum measurement. Entanglement correlations are still difficult to explain with our classical eyes [422].

A 2019 experiment conducted at the University of Glasgow has even allowed to photograph a representation of the state of entangled photons [423]. Some other proposals also exist to see entanglement with your own eyes but has not yet been implemented [424, 425].

Hyperentanglement. You can generate entanglement between different properties or degrees of freedom of a same particle [426, 427], like the polarization and position of a photon, and simple molecules [428]. This is called hyperentanglement [429]. I wouldn't dare trying to figure out what it means with a realistic explanation! The spatially varying polarization correlations in a hyperentangled state can enable quantum holographic imaging or the creation of photonic cluster states for efficient quantum computation with some topological protection effects [430]. Photon polarization and frequency hyperentanglement can also be used in quantum communications [431].

Entanglement with no common past. Nevertheless, we are still able to entangle quantum objects that do not necessarily have a common past [432]. Bell's inequalities were first validated with photons in the visible spectrum. It has been extended to other parts of the spectrum, of course in the infrared bands that are used in fiber optics and free space quantum communications and even in the X-ray band [433, 434]. It has also been done with all sorts of qubits (superconducting, silicon spin, trapped ions, neutral atoms). An EPR experiment was even

realized in 2023 with two times about 700 atoms in a BEC (Bose-Einstein Condensate) [435].

Despite its randomness happening at measurement, entanglement is a very powerful resource. It helps generate random secret keys for two parties with the QKD (quantum key distribution) protocols. It powers quantum computing by creating interdependencies between qubits. Multi-qubits quantum gates conditionally link them together. Once entangled, qubits have inseparable quantum states. Without it, no useful quantum algorithm could work. But quantum entanglement does not mean we can transmit some useful information faster than light since the entangled objects' properties are random.

Bell inequality test. Based on the linearity of quantum mechanics, a Bell inequality test or Bell experiment (see glossary, and Figure 118) looks at the statistical correlation between the states of two quantum objects, with an experiment done a large number of times with the same settings. The test is passed as soon as the statistical correlation S exceeds 2, which is the ceiling imposed by local realism. Quantum mechanics predicts values up to $2\sqrt{2} \approx 2.83$, the Tsirelson bound, which is the maximum correlation reachable by any quantum state and measurement setting, and real experiments land between the two. It demonstrates indirectly that the entangled objects correlated measurement can't be explained by a hidden variable or some form of information broadcasting [436]. Bell's theorem states that no local hidden-variable theory can reproduce all the predictions of quantum mechanics. Put differently, any theory reproducing the observed correlations must give up either locality or the assumption that measurement outcomes are fixed in advance by variables carried by the particles. In photonic experiments, this is obtained with the two photon polarizers having an angle difference of 22.5° and 67.5° as shown in Figure 117. With many other angles, the modulus of S stays below 2 and cannot violate Bell's inequality and the measurement correlations on the two photons detectors are said to be "classical". For Bell inequality tests implemented with regular spin 1/2 qubits like with superconducting qubits defined in their Bloch sphere, the angles maximizing Bell's inequality tests are doubled, at 45° and 135° .

In 2025, a Bell inequality violation experiment was achieved with unentangled photons by Anton Zeilinger, Markus Aspelmeyer and a team of researchers from China and Austria, using a nonlinear crystal based interferometer [437]. How is that possible? The authors say that indistinguishability of the photons creation processes and an active manipulation of photon states during their creation, rather than entanglement, explain this result. But other researchers argue that this result is an artifact of the postselection that is used in the experiment [438]. In a word, the experiment created a "loophole", breaking the Bell's assumption of statistical independence between the photons measurement.




Figure 117: The S value depends on the angle between the two polarizers used. To maximize Bell inequality to $2\sqrt{2}$, you need angles of 22.5° and 67.5° with photon polarizers or 45° and 135° with qubits. Recreated in Python from [source](#).

Loopholes free tests. A loophole-free Bell inequality test is an experiment that aims to distinguish between the predictions of quantum mechanics and those of local realism, by eliminating all known “loopholes” that could otherwise allow a local-realist explanation for observed correlations.

The known loopholes are multiple:

- **Communication loophole** if the information about a measurement setting or outcome on one side of the experiment can travel to the other side and influence the measurement outcome there, even if that communication is slower than light but still faster than the measurement process. To close this loophole, the two measurement stations must be spatially separated by a sufficient distance, and the measurement settings must be chosen randomly - like using a QRNG that we’ll describe later in the book - and quickly, ensuring that there is no time for any signal, even at the speed of light, to travel between the stations before the measurements are completed. This ensures spacelike separation of the measurement events on both sides of the experiment. The distance must be above the speed of light multiplied by the measurement time, which therefore must be as fast as possible. Current experiments use very large distances of many kilometers.
- **Detection loophole** which occurs when the detectors used in the experiment are not perfectly efficient, meaning that only a subset of the emitted entangled particles are well detected. If the detected subset is not a fair or representative sample of all emitted pairs, a local-realist model could potentially explain the observed correlations in the detected sample, even if the entire ensemble of particles would satisfy Bell’s inequality. Closing this loophole requires a detection efficiency above a threshold that depends on the state and the inequality. It is about 82.8% for a maximally entangled state tested with CHSH, and drops to $2/3$ (about 66.7%) for suitably non-maximally entangled

states tested with the Eberhard inequality, which is the route the 2015 photonic experiments took.

- **Freedom-of-Choice Loophole** happens when the choices of measurement settings by the experimenters, or random number generators used for this purpose, are not truly independent but are somehow influenced by hidden variables or past events. It could lead to apparent quantum correlations. This loophole is avoided using truly random and unpredictable methods for setting selection are required, like with a QRNG.


The first loophole-free Bell inequality violation experiments were all conducted in 2015: first by Ronald Hanson’s team at TU Delft in August, using entangled electron spins from nitrogen-vacancy (NV) centers separated by 1.3 km [439], then in November by the team of Anton Zeilinger in Vienna [440] and independently by Lynden Shalm’s team at NIST, both with entangled photons [142]. A follow-up Delft run with additional analysis was published in 2016 [441]. In 2017, an event-ready loophole-free test was achieved at LMU Munich by Wenjamin Rosenfeld, Harald Weinfurter et al., with two rubidium atoms separated by 400 m, closing the detection and locality loopholes simultaneously [442]. Separately and the same year, entanglement was generated between distant quantum dot spins at Cambridge by a team including Robert Stockill, now CTO of QPhoX (the Netherlands), and Claire Le Gall, VP of engineering at Nu Quantum (UK) [443].

Mermin inequalities test created by David Mermin in 1990 extends Bell’s inequalities test to the entanglement of a higher number of quantum objects like a GHZ state with three or more qubits. Daniel Greenberger, Michael A. Horne and Anton Zeilinger imagined a four entangled particles thought experiment in 1989. David Mermin then simplified it to use only three particles, which became the basis of the simplest GHZ state also labelled $|GHZ_3\rangle$ [444]. The record was pushed in 2024 in China with a 60-qubit GHZ state on superconducting qubits [445], then to 75 qubits by Q-CTRL on an IBM Heron r2 QPU on a single-shot basis [446], and to 120 qubits on an IBM QPU with readout error mitigation [447].

These tests are very costly as you increase the number of correlated quantum objects. Another variation is to conduct a state tomography for a set of qubits. Again, its cost grows exponentially with the number of qubits, which explains why most qubit tomographies are not done beyond 6 qubits.

One key aspect of entanglement must be considered: two or more quantum objects are entangled in relation to their measurement context. Two particles that are entangled in one context may be separable in another context. One example is the helium atom, depending on its measurement basis with individual electron spins or with the singlet/triplet basis [391].

realism and locality



$$S = \langle A_0 B_0 \rangle + \langle A_0 B_1 \rangle + \langle A_1 B_0 \rangle - \langle A_1 B_1 \rangle \leq 2$$

average of correlations between A_0, A_1 being equal to +1/-1 and B_0, B_1 being equal to +1/-1, A_n and B_n correspond to binary measurement of two properties 0 and 1, yielding +1 or -1, assuming these properties independently of being measured (realism) and that the choice to measure properties 0 and 1 by Alice cannot influence Bob choice (locality). This is CHSH's inequality.

quantum mechanics

$$\frac{|00\rangle + |11\rangle}{\sqrt{2}} \quad \begin{matrix} A_0 = \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \\ A_1 = \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \end{matrix} \quad \begin{matrix} B_0 = -\frac{\sigma_x + \sigma_z}{\sqrt{2}} \\ B_1 = \frac{\sigma_x - \sigma_z}{\sqrt{2}} \end{matrix} \quad \begin{matrix} \langle A_0 \otimes B_0 \rangle = \langle A_0 \otimes B_1 \rangle = \langle A_1 \otimes B_0 \rangle = \frac{1}{\sqrt{2}} \\ \langle A_1 \otimes B_1 \rangle = -\frac{1}{\sqrt{2}} \end{matrix}$$

Victor prepares a Bell state

Pauli matrices σ_z and σ_x play the role of A_0 and A_1 in the above experiment.

B_0 and B_1 measurements are done in another basis.

$$S = \langle A_0 \otimes B_0 \rangle + \langle A_0 \otimes B_1 \rangle + \langle A_1 \otimes B_0 \rangle - \langle A_1 \otimes B_1 \rangle = 2\sqrt{2}$$

Tsirelson bound

Figure 118: The famous statistical rules computing the test statistics S for two-qubit pure states which separates classical realism from quantum mechanics entanglement, when it matches the Tsirelson bound which has other values for entanglement states with more than one inputs for each quantum object. It is using four different contexts of measurements for two entangled quantum object properties. Source: Wikipedia. 2023.

Polytopes are geometric representation of sets of probability distributions. These are generalizations of polygons and polyhedrons. They represent sets of possible probability distributions that can arise from measurements. Non-contextual polytopes represent the set of probability distributions that would be possible if the system were describable by a non-contextual hidden variable theory. They are generally a subset of a higher-dimensional space and convex, meaning that any linear combination of points within this polytope also lies within it. Contextual polytopes represent the set of all probability distributions that can be achieved by quantum systems. It can have a more complex structure than the non-contextual one, capturing the essence of contextuality. The vertices of this polytope correspond to extreme quantum states and measurements. You can see these polytopes representations in many papers on quantum contextuality [448]. Figure 119 presents a Bell state polytope 2D slice showing the region with quantum correlations in purple.

Parameter and outcome independence. Local realism is based on two assumptions: parameter independence and outcome independence. Parameter independence happens when the outcome of a measurement on one part of a system does not depend on the settings of the measurement apparatus on a spatially separated part. It is a locality condition. Outcome independence on one part of a system does not depend on the outcomes of measurements on a spatially separated part.

Violating parameter independence means that the settings on one side can affect the outcomes on the other side directly. However, quantum mechanics primarily violates the combination of parameter independence and outcome independence, leading to non-local correlations.

Non-locality, as demonstrated by the violation of Bell inequalities by Bell tests ala Alain Aspect, shows that no local hidden variable theory can reproduce the predictions of quantum mechanics. This does not necessarily mean that parameter independence is violated in every quantum mechanical scenario. Instead, quantum mechanics can maintain parameter independence while violating outcome independence, or vice versa, but their combination leads to the overall non-locality that is empirically observed.

Maximally entangled state, aka a Bell state or EPR pair, is a quantum state of a composite system where the subsystems are highly correlated in such a way that the state cannot be factored into tensor products of separate states of its subsystems. The quantum state of the whole system is described entirely in terms of the correlations between the subsystems, with no information about the individual states of the subsystems themselves.

In the context of a bipartite system (consisting of two subsystems), a maximally entangled state can be represented as:

$$|\psi\rangle = \frac{1}{\sqrt{2}} \sum_{i=0}^{d-1} |i\rangle_A \otimes |i\rangle_B$$

where d is the dimension of the Hilbert space of each subsystem, $|i\rangle_A$ and $|i\rangle_B$ are the basis states of the subsystems A and B , respectively. For a two-qubit system with $d=2$, the four maximally entangled Bell states are:




Figure 119: Bell polytope showing correlations between Alice and Bob measurement outcomes and when they are “quantum”. The X and Y axis correspond to Alice and Bob measurement expectation values correlation on two different measurement basis sets. Source: [449]. Added in 2024.

$$\begin{aligned}
 |\Phi^+\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) |\Phi^-\rangle \\
 &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle) \\
 |\Psi^+\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle) |\Psi^-\rangle \\
 &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)
 \end{aligned}$$

They exhibit perfect correlations in the measurement outcomes of the two qubits. Maximally entangled states have several key properties: non-locality (subsequent measurements are correlated whatever the distance), information correlation (the entire information about the state of the composite system is encoded in the correlations between the subsystems) and quantum entanglement measures (it maximizes measures of entanglement, such as the von Neumann entropy of the reduced density matrix).

Non-maximally entangled state: it is a quantum state of a composite system where the subsystems are entangled but not to the maximum extent possible. The quantum correlations between the subsystems are weaker than those in a maximally entangled state. For a bipartite system, non-maximally entangled states can be represented by states that are not fully entangled, like for example:

$$|\psi\rangle = \alpha|00\rangle + \beta|11\rangle \text{ with } \alpha \neq \frac{1}{\sqrt{2}} \text{ and } \beta \neq \frac{1}{\sqrt{2}}$$

The degree of entanglement of the composite system can be quantified either with the von Neumann entropy of the reduced density matrix, or, more simply for a pure two-qubit state, with the **concurrence** C :

$$C = 2|\alpha\beta|$$

For a maximally entangled state, $C = 1$ and for a non-maximally entangled state, $0 < C < 1$. When $C = 0$, we have a basis state with no entanglement at all.

GHZ, W and Dicke states are various entangled states that can be prepared, particularly with gate-based quantum computers.

A GHZ state (Greenberger–Horne–Zeilinger) is a maximally entangled quantum state involving three or more qubits. With 3 qubits, it has the form:

$$|\text{GHZ}_3\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)$$

And with N qubits:

$$|\text{GHZ}_N\rangle = \frac{1}{\sqrt{2}}(|00\dots 0\rangle + |11\dots 1\rangle)$$

With a GHZ state, all the qubits are perfectly correlated and either all are in the $|0\rangle$ state or all are in the $|1\rangle$ state. This state demonstrates strong nonlocal correlations and is often used in tests of quantum foundations and Bell-type inequalities. GHZ states are extremely fragile. If one qubit is lost or measured, the remaining

system loses its entanglement entirely and collapses into a classical mixture. This sensitivity makes GHZ states powerful for certain tasks, but difficult to preserve in noisy or lossy environments.

A W state is another type of multipartite entangled state although not maximally entangled. It distributes a single excitation across multiple qubits in a symmetric way. The three-qubit W state is:

$$|W_3\rangle = \frac{1}{\sqrt{3}}(|100\rangle + |010\rangle + |001\rangle)$$

And the generic N-qubit W state is:

$$|W_N\rangle = \frac{1}{\sqrt{N}}(|100\dots 0\rangle + |010\dots 0\rangle + \dots + |000\dots 1\rangle)$$

Unlike GHZ states, W states are more robust against qubit loss. If one qubit is lost or measured, the remaining qubits still retain entanglement. This makes W states particularly useful for distributed quantum information tasks, such as quantum networking and quantum secret sharing. Although they do not exhibit maximal entanglement like GHZ states, their resilience makes them a practical and valuable resource in multi-party protocols.

Dicke states generalize the idea behind W states by allowing more than one excitation (i.e., more than one qubit in the $|1\rangle$ state) to be symmetrically distributed among the qubits.

For example, a four-qubit Dicke state with two excitations is an equal superposition of all binary strings with two 1s and two 0s:

$$|D_4^2\rangle = \frac{1}{\sqrt{6}}(|1100\rangle + |1010\rangle + \dots + |0101\rangle + |0011\rangle)$$

And the generic N-qubit Dicke state looks like:

$$|D_k^N\rangle = \frac{1}{\sqrt{\binom{N}{k}}} \sum_{\text{perms}} |1\dots 1 0\dots 0\rangle$$

The sum contains all distinct permutations with k excitations among N qubits.

Dicke states are highly symmetric and possess multipartite entanglement that is generally robust against the loss of individual qubits. They are important in quantum metrology, where their collective properties can be used to enhance measurement precision beyond classical limits. Their symmetry and resilience also make them promising for fault-tolerant quantum sensing and computation.

They provide a tunable trade-off between entanglement strength and robustness, especially useful in metrology and error-tolerant protocols.

In summary, GHZ, W, and Dicke states are distinct types of multipartite entangled states, each suited to different quantum applications. GHZ states provide strong nonlocal correlations but are fragile. W states sacrifice some entanglement strength for exceptional robustness. Dicke states offer a tunable balance between the two, depending on the number of excitations and qubits involved, making them valuable in precision sensing and resilient quantum protocols.

Energy-time and time-bin entanglement are two other types of entanglement worth mentioning [450].

Energy-time entanglement was proposed by J. D. Franston in 1989 [451], this form of entanglement arises when two photons are created simultaneously but with uncertain emission times. The photons are entangled so that their energy, defined by their wavelength, and arrival time are correlated. This can be used in quantum key distribution, quantum teleportation, and Bell test experiments to limit the effects of decoherence in optical fibers.

Time-bin entanglement is a refinement of energy-time entanglement, introduced in the late 1990s where qubits are encoded in discrete time bins, typically “early” and “late” arrival times of photons in their detectors [452]. It splits a single photon into two paths using a Mach-Zehnder interferometer (MZI) using one longer path (a “delay line”) that creates a delay. The photon then exists in a superposition of being in the early or late time bin. When two such photons are entangled, their time-bin states are correlated. The entangled photons are still entangled, but with some delay difference. Creating a Bell inequality test experiment in that case must account for the time difference in the photon arrival on the detectors when measuring correlations. This is used for long-distance quantum communication and is more robust than polarization entanglement in fiber optics.

Entanglement monogamy comes from the fact that maximally entangled states of two parties can’t be entangled with other qubits. As a result, in a quantum algorithm with several qubits, all qubits are progressively partially entangled [453].

In science at the frontier of science fiction, some imagine exploiting quantum entanglement to analyze a quantum state inside a black hole [454]! Superposition also happens within benzene C_6H_6 with two carbon-carbon links with their neighbors, using one or two electrons.

Which quantum objects?. Initially, most Bell-inequality tests done to invalidate Bell’s inequalities were realized with pairs of entangled photons generated by single atoms excitation, like calcium in the case of Alain Aspect’s 1982 experiment. Subsequent experiments were also done with photons, particularly in the wake of the development of quantum cryptography, and then many other quantum objects, with atoms, atom ensembles (although not loophole free in that case), trapped ions, electrons and even sometimes, molecules. You can also entangle quantum object of different na-

ture and properties. One example is the entanglement between pairs of free electrons in NV center cavities and the surrounding atom nucleus spins (^{13}C and nitrogen). You can also entangle atoms and photons, and so on.

The most extreme case of entanglement experiment is implemented with elementary particles from the standard model beyond photons and electrons, that are mainly tested at the Large Hadron Collider at CERN in Geneva, like with Z bosons and muons [455, 456].

Entanglement distance. On paper, a Bell inequalities test can be implemented with photons received in two different locations separated by a long distance. With satellite QKD experiments, the record distance is of about 1,200 km. A proposal was made in 2026 to implement this experiment with one polarimeter on the Moon and the other on the Earth or on a satellite positioned at one of the Lagrange points [457], thus avoiding the perturbations and losses generated by the Earth atmosphere. Another older proposal was made using Mars [458].

Entangled mixed states. Mixed states can also be entangled, when they cannot be written as a convex mixture of pure states like $\rho_{AB} = \sum_i p_i \rho_i^A \otimes \rho_i^B$, p_i being classical probabilities and ρ_i^A and ρ_i^B sets of pure states. There is a complex hierarchy of such states with a distinction between entangled mixed states, steering and mixed states violating Bell inequalities, on top of quantum discord (see the hierarchy and references in Figure 120). See 4.1.8 for a primer on pure and mixed states.

3.5 Indetermination

Heisenberg’s principle of indeterminacy or indetermination states that one cannot accurately measure both the position and momentum of a particle or two complementary quantities describing a quantum object state. It is mathematically described as an inequality showing that the multiplication of both precisions cannot be lower than the Planck reduced constant divided by 2:

$$\Delta x \Delta p \geq \hbar/2$$

This inequality was not created by Werner Heisenberg but devised and demonstrated by **Earle Hesse Kennard** in 1927 as he was doing a sabbatical at the University of Göttingen. It is even named the Kennard inequality or Heisenberg-Kennard inequality [468, 469].

The indeterminacy principle has another consequence: one cannot observe at the same time a quantum object in its particle state and in its wave state, per the principle of complementarity enacted by Niels Bohr around 1928, that we already mentioned in the wave-particle duality section. It also explains vacuum quantum fluctuations that we cover later.

For purists, the notions of particle speed and position are even meaningless for electrons. Its characterization is based on its wave nature and its probabilistic description via Schrödinger’s wave function. Don’t even try to understand where it is at a given place and time.

When it deals with momentum and position or waves, Heisenberg’s indeterminacy principle is closely related to a characteristic of Fourier transforms: a nonzero function and its Fourier transform cannot both be sharply concentrated, so, precisely measurable. The more concentrated a signal is in the time domain, the more spread out it is in the frequency domain and vice versa. We have here a mathematical balance between a pulse length precision and its spectral analysis precision (Figure 121).

Since complementary (or incompatible) properties can’t both be measured with an arbitrary precision, we can improve one property measurement precision by decreasing the measurement precision of the complementary property. It is being implemented with the so-called photons squeezing technique. This technique is implemented in the latest LIGO (USA) and VIRGO (Italy) huge interferometers that are used to measure gravitational waves generated by huge astrophysical phenomena like dual black hole collapses. These instruments squeeze the phase quadrature of the light circulating in the interferometer, which lowers the shot noise on the fringe readout, at the price of greater amplitude (photon number) noise that shows up as radiation pressure noise pushing on the mirrors. Frequency-dependent squeezing, deployed on both detectors since 2023, lets each quadrature be squeezed where it matters [470]. Kip Thorne (1940, USA), Rainer Weiss and Barry C. Barish got the Nobel prize in physics in 2017 for their contributions to the creation of the LIGO detector and the observation of gravitational waves.

The indetermination principle is applicable to complementary properties which usually happen to be continuous variables, and not quantified ones like the 4 electron quantum numbers in an atom. One reason is that these 4 properties are compatible. They can theoretically be measured independently and in any order. When a property is quantified, you can’t compute a measurement outcome precision like with continuous variables. The discretized measurement outcome is governed by Max Born’s rule, which is a probabilistic equation (see Postulate IV in 3.1).

Phase space is a framework providing a description of a physical system’s state encapsulating two conjugate variables like position and momentum. In classical mechanics, phase space describes well-defined variables representing the exact position and momentum of particles. With quantum mechanics, inherent uncertainties modify this picture. The phase space is often represented by the Wigner quasi-probability distribution function which assigns a quasi-probability density to each point in phase space, noted using $W(q, p)$. This function can take negative values, showcasing the non-classical nature of quantum mechanics. The quantum phase space




Figure 120: hierarchy of bipartite quantum correlations for pure states (left) and mixed states (right). For a pure state, the distinctions collapse. The quantum state is either a separable product state with no correlations, or it is entangled and non separable. Gisin’s theorem guarantees that every pure entangled state also violates a Bell inequality [459]. Mixed states introduce different kinds of nested states. Classical correlations covers all states, including zero-discord states in yellow [460, 461]. Quantum discord corresponds to mixed states states whose correlations admit no fully classical description [462], a property shared by almost all states, including many separable ones [463]. Entanglement or nonseparability is a subset of the discordant states [464]. Steering corresponds to one party remotely collapsing the other’s conditional states beyond any local-hidden-state model [465]. Bell violation (or Bell nonlocality) is the most demanding property [466]. This is schematic and not a universal set inclusion. © Olivier Ezratty, 2026, inspired by [467].




Figure 121: A simple depiction of a frequency-time spread of a single frequency photon and photon pulses with two different durations and frequency spreads. These explain Heisenberg’s indeterminacy principle over these two complementary observables: the sharper one observable is (time or frequency), the shallower the other will be. The number frequency spread is usually continuous, even with the frequency combs generated by short-pulse lasers, a technique that we describe in the quantum sensing part of this book. © Olivier Ezratty, 2026.

respects the Heisenberg indeterminacy principle, which states that the two conjugate variables cannot both be precisely known simultaneously. On top of position and momentum, other typical conjugate variables are angle and angular momentum (for example, for an electron), time and energy [469], phase and photon number, and

electric and magnetic fields. This principle is inherently encoded in the structure of the Wigner function and in the non-commutative relation $[\hat{q}, \hat{p}] = i\hbar$. Quantum phase space methods are useful in quantum optics, quantum information, and semiclassical approximations. They provide a framework for visualizing and computing

Source and pulse type	Pulse duration (Δt)	Bandwidth (Δf)
Femtosecond laser (Ti:Sapphire)	100 fs (10^{-13} s)	4.41 THz
Quandela quantum dot (single photon) [471]	150 ps (1.5×10^{-10} s)	2.94 GHz
Microwave pulse (5 GHz carrier) [472]	20 ns (2×10^{-8} s)	22.05 MHz

Table 4: Calculated bandwidth (Δf) for standard Gaussian pulse durations (Δt) based on the time-bandwidth product $\Delta t \cdot \Delta f \approx \frac{2 \ln(2)}{\pi} = 0.441$.

quantum phenomena, bridging classical intuition with the abstract formalism of quantum mechanics. The Wigner function is used to study coherence and entanglement in quantum systems, analyze decoherence processes, and develop quantum tomography techniques for reconstructing quantum states from experimental data. The density operator takes the place of the classical phase space coordinates [473].

3.6 Measurement

Measuring quantum object properties follows a very different path than with classical physics due to the back action induced by quantum measurement on the measured system and to its probabilistic dimension.

3.6.1 Continuous properties measurement

With classical mechanics, you can usually predict over time the results of the measurement of macro-objects properties (dimension, speed, position) based on their dynamics. In quantum mechanics, given the knowledge of the position of the measured object, one cannot measure precisely its momentum. More generally, the knowledge we have about two non-commuting observables is bounded as we saw. We cannot assign them a well defined value simultaneously, due to Heisenberg indeterminacy principle. This is the case for the measurement of quantum continuous properties or dynamic variables (in green in Figure 87). With discrete properties like the various electron quantum numbers, measurements can theoretically be done on each properties in any order since they are compatible.

3.6.2 Discrete properties measurement

With discrete properties, before measurement, a single coherent isolated quantum object is said to be in a pure state, represented by a vector in a Hilbert space, or its “psi” (ψ) vector. It is a superposition, or linear combination of basis states or one of the object basis states, like “ground state” or “excited”. When a quantum object is measured against one observable, the state of the quantum object becomes one of the observable basis states, like a spin direction up or down or a discrete energy level. The quantum object “collapses” in a probabilistic way into one of the available basis states

which is detected classically and yields some classical information. When conducting another measurement, one always obtains the same result being the basis state that was obtained beforehand in the first measurement. This is also called “*Schrödinger wave function collapse*” or “*wave packet collapse*” which however works only with so-called projective measurement, as defined by John von Neumann.

With a photon of intermediate polarization between horizontal or vertical linear polarization, it will become a horizontally or vertically polarized photon after its polarization measurement.

Moreover, a quantum measurement readout requires some interaction with a macroscopic object that automatically selects one specific outcome. In straw-man language, quantum measurement is in the eye of the beholder! Measuring the same initial state several times can lead to different outcomes. However, even if each measurement yields a probabilistic result, when repeated several times, their statistical distribution is not probabilistic. It corresponds to the knowledge that can be obtained from the evaluated quantum state created experimentally in a similar way several times.

3.6.3 Qubit measurement

In quantum computing, this principle of reduction is implemented when measuring the state of a qubit. It modifies its value by collapsing it to the basis states $|0\rangle$ or $|1\rangle$.

The outcome is probabilistic with a chance of retrieving a $|0\rangle$ or a $|1\rangle$ depending on the qubit state. However, when the quantum state is a basis state, say $|0\rangle$ or $|1\rangle$ for a qubit, its measurement should return this basis state in 100% of the cases and is therefore not probabilistic but deterministic. This works, however, only in a perfect world without any quantum noise. Even when a qubit is in a basis state, its measurement does not return a perfect basis state 100% all the time. You get a % that is inferior to 100% and corresponds to the readout qubit fidelity. It turns a basis state measurement into a probabilistic one. And, worse, an ideal quantum measurement has an infinite cost, so we will deal with imperfection [474].

The subtle information contained in a qubit that is represented by a complex number or a two-dimensional vector is reduced to $|0\rangle$ or $|1\rangle$ at the time of its measurement. It




Figure 122: Quantum measurement explained with qubits, © Olivier Ezratty, 2021-September 2024.

becomes a classical bit. A single measurement makes us lose all the wealth of information contained in the qubit. We turn the equivalent of two floating point numbers into a single bit. However, this measurement is supposed to happen only at the end of quantum algorithms. During computing, the whole wealth of qubit internal information is leveraged, particularly with the creation of interferences between qubits. All this is illustrated in Figure 122. We will come back to the meaning of α and β complex numbers in the next section on qubits.

This reduction should occur theoretically only at the end of computing. During computing, qubits are modified by quantum gates preserving the richness of their information, the combinatorial nature of their values based on superposition and entanglement. However, quantum measurement is to be implemented during computing with systems implementing quantum error corrections.

The subject of quantum measurement is quite broad. In a forthcoming more detailed section (4.7), we will cover several additional concepts such as projective (von Neumann) measurement, non-selective measurement, weak measurement, gentle measurement and non-destructive measurement.

3.6.4 Quantum speed limits

A quantum speed limit (QSL) is a lower bound on the time a quantum system needs to evolve from one state into a measurably different one. It is also understandable as a dynamical sharpening of the energy-time uncertainty relation [475–477]. It describes the pace at which a state can change in the Hilbert space [478]. It can be derived to find lower bounds of qubit gate [479] and readout speeds [480, 481], considering that actual qubit operation times also depend on operational parameters like control signals bandwidth, noise, operation fidelities and leakage avoidance needs.

QSL is described by two bounds which define the time τ required to reach an orthogonal state for unitary evolutions under a time-independent Hamiltonian.

- The **Mandelstam-Tamm bound** set by the energy spread [483], where ΔE is the standard deviation of the Hamiltonian:

$$\tau \geq \frac{\pi \hbar}{2 \Delta E}$$

- The **Margolus-Levitin bound** set by the mean energy above the ground state [484]:

$$\tau \geq \frac{\pi \hbar}{2 (\langle E \rangle - E_0)}.$$

The physically correct bound is the larger of the two bounds [485]:

$$\tau_{\text{QSL}} = \max \left\{ \frac{\pi \hbar}{2 \Delta E}, \frac{\pi \hbar}{2 (\langle E \rangle - E_0)} \right\}$$

The above bound is applicable to unitary evolution, whereas measurement is a non-unitary operation. Applying a speed limit to measurement requires a modeling choice, one considering measurement as being a globally unitary system and apparatus interaction followed by decoherence. It bounds the time for the joint state to reach the post-measurement state. The cost of a fast measurement is then paid in the measurement device spent energy [486].

In quantum computing, qubit measurement time is always several orders of magnitude slower than the related QSL due to practical reasons across most qubit modalities as shown in Figure 123.

3.6.5 The quantum measurement problem

Understanding how quantum measurements operate at the physical level remains an open question with lots of interpretations [487]. It is folded into the notion of the “measurement problem”, which is still debated among quantum physicists [488], including with David Mer-




Figure 123: With quantum computing hardware, the QSL is almost never the operational limit. Qubit readout is bounded by signal-to-noise, cavity linewidth, dispersive shift, and amplifier noise, which are far above the energy-time floor. For example, a superconducting transmon driven at $\omega/2\pi \approx 5$ GHz prepared in an equal superposition has $\Delta E = \hbar\omega/2$, so $\tau_{\text{QSL}} = \pi/\omega \approx 100$ ps. Real dispersive readout takes 50 to 1000 ns, so devices operate three to four orders of magnitude above the QSL fundamental floor. For trapped-ion fluorescence readout, the gap is larger since state detection takes 100 μs to a few ms, a value set by photon scattering and collection efficiency rather than by any energy-time bound. The gap is even larger with photon readout based on a SNSPD detector, which we describe later in the book, in the photon qubit part. A 2026 paper did show that it was possible to reach gate times of tens of femtoseconds (10^{-14} s), below computed bound conditioned by the Mandelstam-Tamm bound shown in the chart [482]. The reason is that the chart bound is related to the use of microwave or narrow-band continuous-wave lasers to drive hyperfine or Rydberg transitions while the experiment is using ultra-fast femtosecond laser pulses. © Olivier Ezratty, September 2026.

min according to which there is no such measurement problem [489].

A common one explains that a quantum measurement happens when the quantum system being measured is entangled with the measurement apparatus [490]. The measurement outcome is a partial trace of the density matrix associating these two elements.

Another one is the **Koopman-von Neumann** (KvN), a description of classical systems and measurements [491]. It can also use ensembles on configuration spaces describing how a classical system measuring the position of a quantum particle prepared in a superposition of two localized states operates [492, 493].

One disadvantage of many of these representations is that they are not easily applicable to all quantum systems like when you measure a superconducting qubit state with microwave reflectometry or the polarization of a photon. Another one could be the reliance on type III von Neumann algebra with infinite operators, which could describe how the classical-quantum frontier operates [494].

3.7 No-cloning

The no-cloning theorem prohibits the identical copy of an unknown state of a quantum object onto another quantum object. The theorem is mathematically demonstrated in [six lines](#)^w. The theorem was demonstrated in 1982, independently, by William Wootters and Wojciech Zurek [495] and by Dennis Dieks. The article is still not available in open-source on a site such as arXiv,

self-applying the no-cloning theorem [496]! See below a simple demonstration of the theorem.

As a result from the no-cloning theorem, it is impossible to copy the state of an unknown qubit to exploit it independently of its original, contrary to a classical bit that can be copied from/to memory and from/to storage. It also prevents quantum computers from implementing a von Neumann computing model with separate processing and memory (Figure 124).




Figure 124: No-cloning theorem practical consequences for quantum computing. You can't copy some qubit data independently but you can entangle it with other qubits or teleport it to other qubits. © Olivier Ezratty, 2023-2026.

In quantum computers, qubits can be duplicated via quantum gates and entanglement, but the resulting

qubits are entangled and therefore somehow synchronized and inseparable. Reading the copy destroys the original, since measuring one of two entangled qubits projects both onto a basis state, the outcome being drawn at random according to the Born rule. This has a direct impact on the design of quantum algorithms and on the error correction codes of quantum computers. These error-correction codes use the trick of projective measurement on a different computational basis as we'll see later.

A derivative of no-cloning is non-deleting. In the case of a qubit, it means it is impossible to reset a qubit from an entangled set of two qubits $|\psi\rangle$, meaning to transform $|\psi\rangle|\psi\rangle$ into $|\psi\rangle|0\rangle$.

Another option exists which is to teleport a quantum object state to another distant quantum object. The initial quantum object will then be reset.

One consequence of the no-cloning theorem is the impossibility to broadcast an unknown quantum state to multiple parties [497]. If you can't do it once, it makes sense that you can't do it several times. In plain terms, when dealing with a quantum state prepared with a quantum circuit containing quantum gates, you can't broadcast to multiple other quantum computers the quantum state in your register.

One thing that can be done is to send the quantum circuit that was used to prepare this state and let the remote parties execute it and sample it, to yield a statistical distribution of interest. It is a bit like if I couldn't broadcast, or publish, this book in a PDF file and had to send all its $\text{\LaTeX}$ components to all readers and ask them to compile it to generate their own PDF file.

3.8 Tunnel effect

The wave-particle nature of matter allows it to cross physical barriers also known as energy walls in some circumstances, depending on the wall thickness and quantum object wavelength. The transmitted wave is usually attenuated after crossing the barrier and its strength depends on the wavelength with regards to the barrier length and composition.

This phenomenon was first accidentally unveiled by **Henri Becquerel** in 1896 when he discovered radioactivity. It did show up with uranium salts decaying, producing alpha rays comprised of two neutrons and two protons (Figure 125). This phenomenon was later explained thanks to quantum physics and wave-particle duality by **George Gamow** (1904-1968, Ukrainian-Russian-American) in 1928.

Just before in 1927, the German physicist **Friedrich Hund** (1896-1997, German) created the formalism explaining electron based tunneling effect.

The tunnel effect is used in superconducting Josephson junctions and exploited in D-Wave quantum annealers where it is used to drive a system of spin qubits

("Hamiltonian", with a given total energy level) towards an energy minimum corresponding to the resolution of a complex combinatorial problem or a search for energy minimum as in quantum matter physics.

The tunnel effect is not the operating principle of conventional logic transistors, though it is exploited elsewhere in microelectronics, in flash memory where Fowler-Nordheim tunneling programs and erases the floating gate, and in tunnel FETs (TFETs) which are researched for sub-60 mV/decade switching. Most transistors make use of the field effect which was patented by Julius Edgar Lilienfeld in 1926. It is implemented in MOSFETs (metal-oxide-semiconductor field effect transistor) and in CMOS (complementary-metal-oxide-semiconductor, that use variants of MOSFETs). These transistors use a metal gate deposited on a silicon-dioxide (SiO_2) and now a "high-K dielectric" as the gate dimension is decreasing with higher densities, to reduce the tunnel effect. The gate voltage determines the transistor conductivity.




Figure 125: Radioactive alpha decay across the Coulomb barrier which occurs when a nucleus emits an alpha particle which is a helium-4 nucleus composed of two protons and two neutrons, that is trapped by the Coulomb barrier due to electrostatic repulsion. Quantum tunneling allows the alpha particle to escape despite lacking the classical energy to overcome the barrier.

3.9 Symmetries

Symmetries play a key role in quantum science, shaping both our understanding of quantum systems and the design of quantum technologies [498]. A symmetry is a transformation such as a rotation, translation, or reflection that leaves certain properties of a system unchanged (invariant) or predictably altered (equivariant). These transformations reveal deep structural regularities in the laws of physics.

Spatial symmetries and geometric structure. In quantum systems, symmetries often emerge from the geometric nature of the objects being studied. For example, molecules in molecular dynamics are modeled as graphs embedded in three-dimensional space. When a

Q Demonstrating the no-cloning theorem

No Cloning Assume we have a unitary operator U_{cl} and two quantum states $|\phi\rangle$ and $|\psi\rangle$ which U_{cl} clones, i.e.:

$$\begin{aligned} |\phi\rangle \otimes |0\rangle &\xrightarrow{U_{\text{cl}}} |\phi\rangle \otimes |\phi\rangle \\ |\psi\rangle \otimes |0\rangle &\xrightarrow{U_{\text{cl}}} |\psi\rangle \otimes |\psi\rangle \end{aligned}$$

Then $\langle\phi|\psi\rangle$ is 0 or 1.

Proof 1:

$$\begin{aligned} \langle\phi|\psi\rangle &= \langle\phi|\psi\rangle \langle\phi|\psi\rangle \langle 0|0\rangle \langle 0|0\rangle \\ &= \langle\phi \otimes 0|\psi \otimes 0\rangle \langle\phi \otimes 0|\psi \otimes 0\rangle \\ &= \langle\phi \otimes \phi|\psi \otimes \psi\rangle \langle\phi \otimes \phi|\psi \otimes \psi\rangle \\ &= (\langle\phi|\psi\rangle)^2 \end{aligned}$$

In the second equality we used the fact that U_{cl} , being unitary, preserves inner products.

Proof 2: Suppose there exists a unitary operator U_{cl} that can indeed clone an unknown quantum state $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$. Then:

$$\begin{aligned} |\phi\rangle \otimes |0\rangle &\xrightarrow{U_{\text{cl}}} |\phi\rangle \otimes |\phi\rangle = (\alpha|0\rangle + \beta|1\rangle) \otimes (\alpha|0\rangle + \beta|1\rangle) \\ &= \alpha^2|00\rangle + \beta\alpha|10\rangle + \alpha\beta|01\rangle + \beta^2|11\rangle \end{aligned}$$

But now if we use U_{cl} to clone the expansion of $|\phi\rangle$, we arrive at a different state:

$$(\alpha|0\rangle + \beta|1\rangle) \otimes |0\rangle \xrightarrow{U_{\text{cl}}} \alpha|00\rangle + \beta|11\rangle$$

This theorem deals with unknown quantum objects. But if you know how to create a given quantum state with a series of gates, you can replicate this sequence multiple times to recreate the target state. Either in parallel with different qubits, or in a serialized way with the same qubits. Practically, due to quantum noise, the resulting states will probably be slightly different. Here there are no cross terms. Thus we have a contradiction and therefore there cannot exist such a unitary operator U_{cl} .

molecule is rotated or translated, its intrinsic properties such as energy levels or interaction potentials typically remain unchanged. This invariance under spatial transformations is known as rotational or translational symmetry.

Symmetries in quantum algorithms and simulations. Such symmetries are powerful tools in quantum algorithm and simulation design. Exploiting them can help reduce computational complexity, build efficient circuits, and improve error resilience [499–503]. More broadly, symmetries help identify conserved quantities, classify quantum phases of matter, and inspire new architectures for quantum devices. The link between the two is **Noether's theorem**, established by **Emmy Noether** in 1918, according to which every continuous symmetry of a system's action corresponds to a conserved quantity: time translation invariance gives energy conservation, space translation gives momentum, and rotation gives angular momentum. Its quantum counterpart is that an observable commuting with the Hamiltonian is a constant of the motion, which is exactly the property exploited when a quantum simulation is restricted to a fixed particle-number or fixed total-spin sector to shrink the Hilbert space it has to explore [504].

Symmetries and quantum error correction. Symmetries also intersect with quantum error correction. The Eastin-Knill theorem shows that continuous symme-

tries such as arbitrary phase rotations cannot be exactly implemented transversally in quantum error-correcting codes. This imposes fundamental limitations on codes that aim to preserve such symmetries while correcting errors [505, 506]. Nonetheless, various error-correcting codes leverage discrete or approximate symmetries to enhance fault tolerance and operational efficiency [506, 507].

Time symmetries in quantum physics. Time symmetries offer another layer of insight. Time-reversal symmetry refers to the invariance of quantum laws under reversal of time's direction. Mathematically, this is expressed through a time-reversal operator acting on quantum states. If a system is time-reversal symmetric, its evolution appears the same whether time flows forward or backward, though this does not imply literal time travel. This symmetry is often broken by interactions such as magnetic fields or spin-orbit coupling. Time-reversal symmetry plays a crucial role in topological quantum systems, such as topological insulators, where it protects edge states from scattering and decoherence. It also appears in quantum thermodynamics, where breaking time symmetry is linked to entropy production and irreversibility [508]. In quantum computing, time symmetry connects to the reversibility of quantum gates, a key feature in designing coherent and potentially energy-efficient operations, modulo the cost of qubit control.



3.9.1 Bibliography and notes

These are the bibliographical references and notes for the **Quantum physics basics** subsection.

- [125] [Experimental realization of Wheeler's delayed-choice GedankenExperiment](#), by Vincent Jacques, E. Wu, Frédéric Grosshans, François Treussart, Philippe Grangier, Alain Aspect, and Jean-François Roch, arXiv, October 2006 (9 pages).
- [142] [A strong loophole-free test of local realism](#), by Lynden K. Shalm, Emanuel Knill, Sae Woo Nam, et al., Physical Review Letters, December 2015 (10 pages).
- [341] [Neutrinos as qubits and qutrits](#), by Abhishek Kumar Jha, Akshay Chatla, and Bindu A. Bambah, arXiv, January 2024 (33 pages).
- [342] [Bohr's Complementarity and Kant's Epistemology](#), by Michel Bitbol et al., 2013 (22 pages).
- [343] [Mathematical Foundations of Quantum Mechanics: An Advanced Short Course](#), by Valter Moretti, arXiv, July 2016 (103 pages).
- [344] [Axiomatic quantum mechanics: Necessity and benefits for the physics studies](#), by J. Jeknic-Dugic, M. Arsenijevic, and M. Dugic, arXiv, November 2017 (23 pages).
- [345] [Lecture Notes for Ph219/CS219: Quantum Information Chapter 2](#), by John Preskill, July 2015 (53 pages).
- [346] [The Hamilton-Jacobi Equation: an intuitive approach](#), by Bahram Houchmandzadeh, arXiv, October 2019 (7 pages).
- [347] [Magnetic clock for a harmonic oscillator](#), by Alessandro Coppo et al., PRA, May 2024 (10 pages)  (preprint on [inspirehep.net](#)).
- [348] [A Spin-Based Pathway to Testing the Quantum Nature of Gravity](#), by Sougato Bose, Gerard J. Milburn, Gavin W. Morley, et al., arXiv, September 2025 (18 pages).
- [349] [What's everything made of?](#), by Charles T. Sebens, American Scientist, October 2019.
- [350] [The Fundamentality of Fields](#), by Charles T. Sebens, arXiv, September 2022 (32 pages).
- [351] [High-precision measurement of the atomic mass of the electron](#), by S. Sturm, F. Köhler, J. Zatorski, A. Wagner, Z. Harman, G. Werth, W. Quint, C. H. Keitel, and K. Blaum, Nature, February 2014  (preprint on [arXiv](#)).
- [352] [Particle and Nuclear Physics Handout #3](#), by Tina Potter, 2022 (124 slides).
- [353] [Introduction to Nuclear Magnetic Resonance Spectroscopy](#), by Dean L. Olson, 2007 (40 slides).
- [354] [In situ electron paramagnetic resonance spectroscopy using single nanodiamond sensors](#), by Zhuoyang Qin, Jiangfeng Du, et al., Nature Communications, October 2023 (8 pages).
- [355] [A Hypothetical Effect of the Maxwell–Proca Electromagnetic Stresses on Galaxy Rotation Curves](#), by D. D. Ryutov, Dmitry Budker, and V. V. Flambaum, The Astrophysical Journal, February 2019 (13 pages).
- [356] [Quantum advantage using high-dimensional twisted photons as quantum finite automata](#), by Stephen Z. D. Plachta, Markus Hiekkamäki, Abuzer Yakaryilmaz, and Robert Fickler, arXiv, June 2022 (20 pages).
- [357] [Photon Diameter Measurements](#), by Geoffrey Hunter, Marian Kowalski, and Camil Alexandrescu, The Nature of Light: What Is a Photon? (AIP Conference Proceedings), Springer, 2003 (10 pages) .
- [358] [Electromagnetic fields, size, and copy of a single photon](#), by Shan-Liang Liu, arXiv, May 2018 (4 pages).
- [359] [The Size and Shape of a Single Photon](#), by Zhenglong Xu, 2021 (22 pages).
- [360] [Limits for realizing single photons](#), by Jan Gulla, Kai Ryen, and Johannes Skaar, arXiv, November 2023 (22 pages).
- [361] [The concept of modes in optics and photonics](#), by René Dändliker, 1999 (6 pages).
- [362] [Roadmap on multimode photonics](#), by Ilaria Cristiani, 2022 (39 pages).
- [363] [Modes and states in quantum optics](#), by Claude Fabre and Nicolas Treps, Reviews of Modern Physics, September 2020  (preprint on [arXiv](#)).
- [364] [Orbital angular momentum of light and the transformation of Laguerre-Gaussian laser modes](#), by L. Allen, M. W. Beijersbergen, R. J. C. Spreeuw, and J. P. Woerdman, Physical Review A, June 1992 (5 pages) .
- [365] [Measurement of Subpicosecond Time Intervals between Two Photons by Interference](#), by C. K. Hong, Z. Y. Ou, and L. Mandel, Physical Review Letters, November 1987 (4 pages) .
- [366] [Generation of non-classical light in a photon-number superposition](#), by J. C. Loredó, C. Antón, B. Reznichenko, P. Hilaire, A. Harouri, C. Millet, H. Ollivier, N. Somaschi, L. De Santis, A. Lemaitre, I. Sagnes, L. Lanco, A. Auffèves, O. Krebs, and P. Senellart, arXiv, November 2018 (8 pages).
- [367] [Generating superposition of up-to three photons for continuous variable quantum information processing](#), by Mitsuyoshi Yukawa, Kazunori Miyata, Takahiro Mizuta, Hidehiro Yonezawa, Petr Marek, Radim Filip, and Akira Furusawa, arXiv, March 2013 (5 pages).
- [368] [Photon-number entanglement generated by sequential excitation of a two-level atom](#), by Stephen C. Wein, Maria Maffei, Paul Hilaire, Niccolò Somaschi, Loïc Lanco, Alexia Auffèves, Pascale Senellart, Carlos Antón-Solanas, et al., Nature Photonics, April 2022  (preprint on [arXiv](#)).
- [369] [Teleportation of a genuine single-rail vacuum-one-photon qubit generated via a quantum dot source](#), by Beatrice Polacchi, Fabio Sciarrino, et al., arXiv, December 2024 (10 pages).
- [370] [Orbital angular momentum of light and the transformation of Laguerre-Gaussian laser modes](#), by L. Allen, M. W. Beijersbergen, R. J. C. Spreeuw, and J. P. Woerdman, Physical Review A, June 1992 .
- [371] [Optical Spin-to-Orbital Angular Momentum Conversion in Inhomogeneous Anisotropic Media](#), by L. Marrucci, C. Manzo, and D. Paparo, Physical Review Letters, April 2006 (11 pages)  (preprint on [arXiv](#)).
- [372] [Q-plate technology: a progress review \[Invited\]](#), by Andrea Rubano, Filippo Cardano, Bruno Piccirillo, and Lorenzo Marrucci, JOSA B, January 2019 (18 pages) .
- [373] [Measuring the Orbital Angular Momentum of a Single Photon](#), by Jonathan Leach, Miles J. Padgett, Stephen M. Barnett, Sonja Franke-Arnold, and Johannes Courtial, Physical Review Letters, June 2002 .
- [374] [OAM-mode sorting with a wavefront twister](#), by Suman Karan, Swati Chaudhary, Harshwardhan Wanare, and Anand Kumar Jha, arXiv, April 2026 (6 pages).
- [375] [Efficient separation of the orbital angular momentum eigenstates of light](#), by Mohammad Mirhosseini, Mehul Malik, Zhimin Shi, and Robert W. Boyd, Nature Communications, November 2013.
- [376] [Topological robustness of orbital angular momentum entanglement in stochastic channels](#), by Tatjana Kleine, Pedro Ornelas, Cade Peters, Zhenyu Guo, Bereneice Sephton, Isaac Nape, Yijie Shen, and Andrew Forbes, arXiv, March 2026 (13 pages).

- [377] [Photon localization: a comparative study](#), by Laura M. Ferguson, Ricardo Rojas Castellano, Kayla Oltman, Joshua G. Fenwick, and Rainer Dick, arXiv, August 2026 (10 pages).
- [378] [Hologram of a Single Photon](#), by Radosław Chrapkiewicz, Michał Jachura, Konrad Banaszek, and Wojciech Wasilewski, arXiv, January 2016.
- [379] [Lecture 1. Basic concepts of statistical optics](#), by mchekho, Max Planck Institute for the Science of Light, 2017 (7 pages).
- [380] [Photon Antibunching in Resonance Fluorescence](#), by H. J. Kimble, Physical Review Letters, September 1977 📄.
- [381] [High-performance semiconductor quantum-dot single-photon sources](#), by Pascale Senellart, Glenn Solomon, and Andrew G. White, 2017 (14 pages) 📄.
- [382] [Utilising the quantum properties of photons and deep learning to improve parameter estimation](#), by Lewis Howard, University of Queensland, December 2022 (130 pages).
- [383] [Generation of squeezed states by parametric down conversion](#), by Ling-An Wu, 1986 (4 pages).
- [384] [Conversion of Gaussian states to non-Gaussian states using photon-number-resolving detectors](#), by Daiqin Su, Casey R. Myers, and Krishna Kumar Sabapathy, arXiv, April 2019 (37 pages).
- [385] [Non-Gaussian Quantum States and Where to Find Them](#), by Mattia Walschaers, arXiv, July 2021 (61 pages).
- [386] [Production and applications of non-Gaussian quantum states of light](#), by A. I. Lvovsky, Philippe Grangier, Alexei Ourjoumtsev, Valentina Parigi, Masahide Sasaki, and Rosa Tualle-Brouiri, arXiv, June 2020 (50 pages).
- [387] [Recent advances in Wigner function approaches](#), by Josef Weinbub and David K. Ferry, Applied Physics Reviews, December 2018 (34 pages).
- [388] [PhD thesis: Modes, States, and Symmetries in quantum Optics for quantum Information and Metrology](#), by Éloi Descamps, arXiv, July 2026 (508 pages).
- [389] [Deterministic preparation of optical squeezed cat and Gottesman-Kitaev-Preskill states](#), by Matthew S. Winnel, Joshua J. Guanzon, Deepesh Singh, and Timothy C. Ralph, arXiv, November 2023 (19 pages).
- [390] [Realization of a crosstalk-free two-ion node for long-distance quantum networking](#), by P.-C. Lai, L.-M. Duan, et al., arXiv, March 2025 (7 pages).
- [391] [The two-spin enigma: from the helium atom to quantum ontology](#), by Philippe Grangier, Alexia Auffèves, Nayla Farouki, Mathias Van Den Bossche, and Olivier Ezratty, Entropy, November 2024 (13 pages).
- [392] [Spin-photon interface and spin-controlled photon switching in a nanobeam waveguide](#), by Alisa Javadi, Peter Lodahl, et al., arXiv, September 2017 (8 pages).
- [393] [A Spin-Photon Interface Using Charge-Tunable Quantum Dots Strongly Coupled to a Cavity](#), by Zhouchen Luo, September 2019 (19 pages) 📄 (preprint on arXiv).
- [394] [The double-slit experiment with single electrons](#), by John Steeds et al., 2023 (2 pages).
- [395] [Single- and two-photon wave-particle superpositions: Theory and experiment](#), by Qing-Feng Xue, Zhong-Xiao Man, et al., PRA, August 2023 📄.
- [396] [Young's Double-Slit Interference Demonstration with Single Photons](#), by Bill J. Luo et al), arXiv, January 2024 (22 pages).
- [397] [Overview: Dynamics with Quantum Trajectories](#), by Robert E. Wyatt, 2008 (24 slides).
- [398] [An introduction to monitored quantum systems and quantum trajectories: spectrum, typicality, and phases](#), by Ryusuke Hamazaki, Ken Mochizuki, Hisanori Oshima, and Yohei Fuji, arXiv, March 2026 (162 pages).
- [399] [To catch and reverse a quantum jump mid-flight](#), by Z. K. Mineev, Michel H. Devoret, et al., Nature, June 2019 📄 (preprint on arXiv).
- [400] [Observation of Quantum Jumps](#), by Th. Sauter, Physical Review Letters, October 1986 📄.
- [401] [A generalized multipath delayed-choice experiment on a large-scale quantum nanophotonic chip](#), by Xiaojiong Chen, Jianwei Wang, et al., Nature Communications, May 2021 (10 pages).
- [402] [Delayed choice experiments: An analysis in forward time](#), by Marijn Waaijer and Jan van Neerven, arXiv, February 2024 (21 pages).
- [403] [Experimental delayed-choice entanglement swapping](#), by Xiao-song Ma, Anton Zeilinger, et al., Nature Physics, April 2012 📄 (preprint on arXiv).
- [404] [Wave-particle duality of C60 molecules](#), by Markus Arndt, Olaf Nairz, Julian Vos-Andreae, Claudia Keller, Gerbrand van der Zouw, and Anton Zeilinger, Nature, October 1999 (3 pages) 📄 (preprint on ResearchGate).
- [405] [Quantum interference experiments with large molecules](#), by Olaf Nairz et al., 2002 (8 pages).
- [406] [Real-time single-molecule imaging of quantum interference](#), by Thomas Juffmann, arXiv, February 2014 (16 pages).
- [407] [Matter-wave interference of a native polypeptide](#), by Armin Shayeghi, Markus Arndt, et al., arXiv, October 2019 (10 pages).
- [408] [How Big Can the Quantum World Be? Physicists Probe the Limits](#), by Philip Ball, July 2021.
- [409] [Real-time optimal quantum control of mechanical motion at room temperature](#), by Lorenzo Magrini, Markus Aspelmeyer, et al., Nature, July 2021 📄 (preprint on arXiv).
- [410] [Quantum control of a nanoparticle optically levitated in cryogenic free space](#), by Felix Tebbenjohanns, Lukas Novotny, et al., arXiv, April 2021 (26 pages).
- [411] [Experimental challenges for high-mass matter-wave interference with nanoparticles](#), by Sebastian Pedalino, Markus Arndt, et al., arXiv, January 2023 (10 pages).
- [412] [Highly fluorinated model compounds for matter-wave interferometry](#), by Jens Tüxen, 2012 (242 pages).
- [413] [Experimental Evidence for a Photon Anticorrelation Effect on a Beam Splitter: A New Light on Single-Photon Interferences](#), by Philippe Grangier, Gérard Roger, and Alain Aspect, Europhysics Letters, 1986.
- [414] [Quantum Superposition of Two Temperatures](#), by Arun Kumar Pati and Avijit Misra, arXiv, December 2021 (7 pages).
- [415] [Quantum superposition of thermodynamic evolutions with opposing time's arrows](#), by Giulia Rubino, Gonzalo Manzano, and Časlav Brukner, Communications Physics, November 2021 (10 pages).
- [416] [Gravitationally-induced wave function collapse time for molecules](#), by Anderson A. Tomaz et al., Physical Chemistry Chemical Physics, June 2024 (14 pages).
- [417] [Macroscopic quantum coherence and entanglement in mechanical systems](#), by Benjamin Yadin and Matteo Fadel, arXiv, March 2025 (23 pages).
- [418] [The Hydrogen Atom as an Entangled Electron-Proton System](#), by Paolo Tommasini, Eddy Timmermans, and A.F.R. de Toledo Piza, arXiv, October 1997 (17 pages).
- [419] [Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell's Inequalities](#), by Alain Aspect, Philippe Grangier, and Gérard Roger, Physical Review Letters, July 1982 📄.
- [420] [How Do You Create Quantum Entanglement?](#), by Chad Orzel, February 2017.
- [421] [Einstein-Podolsky-Rosen paradox in a hybrid bipartite system](#), by Michał Dąbrowski et al., Optica, 2017 (4 pages).
- [422] [Quantum Nonlocality: how does Nature do it?](#), by Marian Kupczynski, arXiv, February 2024 (17 pages).
- [423] [Scientists unveil the first-ever image of quantum entanglement](#), by Paul-Antoine Moreau, July 2019.

- [424] [What does it take to see entanglement?](#), by Valentina Caprara Vivoli, Pavel Sekatski, and Nicolas Sangouard, arXiv, February 2016 (7 pages).
- [425] [Proposal for witnessing non-classical light with the human eye](#), by A. Dodel, A. Mayinda, E. Oudot, A. Martin, P. Sekatski, J.-D. Bancal, and N. Sangouard, arXiv, April 2017 (9 pages).
- [426] [Generation of intraparticle quantum correlations in amplitude damping channel and its robustness](#), by Animesh Sinha Roy, Urbasi Sinha, et al., arXiv, July 2024 (19 pages).
- [427] [Entanglement between dependent degrees of freedom: Quasiparticle correlations](#), by Franziska Barkhausen, Jan Sperling, et al., arXiv, March 2025 (9 pages).
- [428] [On-Demand Entanglement of Molecules in a Reconfigurable Optical Tweezer Array](#), by Connor M. Holland, Yukai Lu, and Lawrence W. Cheuk, arXiv, October 2022 (15 pages).
- [429] [Direct generation of multi-photon hyperentanglement](#), by Peng Zhao, Yu-Bo Sheng, et al., arXiv, June 2024 (16 pages).
- [430] [Rapid comprehensive characterization of biphoton spatial-polarization hyperentanglement](#), by Cheng Li, Robert Boyd, et al., arXiv, February 2025 (16 pages).
- [431] [Photonic hyperentanglement in polarisation and frequency via joint spectrum shaping](#), by Tommaso Faleo, Fabian Steinhauser, Gregor Weihs, Stefan Frick, and Robert Keil, arXiv, March 2026 (16 pages).
- [432] [Qubits that never interact could exhibit past-future entanglement](#), by Lisa Zyga, July 2012.
- [433] [Entangled X-ray Photon Pair Generation by Free Electron Lasers](#), by Linfeng Zhang et al., arXiv, February 2023 (15 pages).
- [434] [Production of entangled x rays through nonlinear double Compton scattering](#), by T. D. C. de Vos, J. J. Postema, B. H. Schaap, A. Di Piazza, and O. J. Luiten, arXiv, October 2024 (17 pages).
- [435] [Einstein-Podolsky-Rosen Experiment with Two Bose-Einstein Condensates](#), by Paolo Colciaghi et al., PRX, May 2023 (10 pages).
- [436] [No-broadcasting characterizes operational contextuality](#), by Pauli Jokinen, Mirjam Weilenmann, Martin Plávala, Juha-Pekka Pellonpää, Jukka Kiukas, and Roope Uola, arXiv, June 2024 (12 pages).
- [437] [Violation of Bell Inequality with Unentangled Photons](#), by Kai Wang, Anton Zeilinger, Xiao-Song Ma, et al., arXiv, July 2025 (13 pages).
- [438] [Bell Inequality Violations Without Entanglement? It's Just Postselection](#), by Ken Wharton and Huw Price, arXiv, August 2025 (8 pages).
- [439] [Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres](#), by B. Hensen, Ronald Hanson, et al., Nature, October 2015 (8 pages).
- [440] [Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons](#), by Marissa Giustina et al., Physical Review Letters, December 2015 (7 pages).
- [441] [Loophole-free Bell test using electron spins in diamond: second experiment and additional analysis](#), by B. Hensen, Anaïs Dréau, Stephanie Wehner, and Ronald Hanson, Scientific Reports, August 2016.
- [442] [Event-Ready Bell Test using Entangled Atoms Simultaneously Closing Detection and Locality Loopholes](#), by Wenjamin Rosenfeld, Daniel Burchardt, Robert Garthoff, Kai Redeker, Norbert Ortegel, Markus Rau, and Harald Weinfurter, Physical Review Letters, July 2017 (6 pages).
- [443] [Phase-Tuned Entangled State Generation between Distant Spin Qubits](#), by Robert Stockill, Claire Le Gall, et al., Physical Review Letters, July 2017 (preprint on arXiv).
- [444] [Extreme quantum entanglement in a superposition of macroscopically distinct states](#), by N. David Mermin, Physical Review Letters, October 1990 (4 pages).
- [445] [Creating and controlling global Greenberger-Horne-Zeilinger entanglement on quantum processors](#), by Zehang Bao, H. Wang, et al., Nature Communications, October 2024 (7 pages) (preprint on doi.org).
- [446] [Achieving computational gains with quantum error-correction primitives: Generation of long-range entanglement enhanced by error detection](#), by Haoran Liao, Michael J. Biercuk, Yuval Baum, et al., arXiv, May 2025 (8 pages).
- [447] [Big cats: entanglement in 120 qubits and beyond](#), by Ali Javadi-Abhari, Simon Martiel, Alireza Seif, Maika Takita, and Ken X. Wei, arXiv, October 2025 (9 pages).
- [448] [Lifting noncontextuality inequalities](#), by Raman Choudhary, Rui Soares Barbosa, and Adán Cabello, arXiv, May 2024 (20 pages).
- [449] [Tsirelson Polytopes and randomness generation](#), by Peter Bierhorst et al., 2020 (20 pages).
- [450] [Energy-time and time-bin entanglement: past, present and future](#), by Guilherme B. Xavier, Adán Cabello, et al., arXiv, April 2025 (27 pages).
- [451] [Bell inequality for position and time](#), by J. D. Franson, Physical Review Letters, May 1989.
- [452] [Observable and computable entanglement in time](#), by Alexey Milekhin, Zofia Adamska, and John Preskill, arXiv, February 2025 (13 pages).
- [453] [Introduction to quantum entanglement in many-body systems](#), by Anubhav Kumar Srivastava, Maciej Lewenstein, Grzegorz Rajchel-Mieldzioć, et al., arXiv, October 2024 (61 pages).
- [454] [Can entangled qubits be used to probe black holes?](#), by Robert Sanders, 2019.
- [455] [About testing Bell locality at colliders](#), by M. Fabbrichesi, R. Floreanini, and L. Marzola, arXiv, March 2025 (19 pages).
- [456] [CERN QT4HEP, LHC, Atlas, CMS and antimatter factory](#), by Olivier Ezratty, Opinions Libres, January 2025.
- [457] [Quantum correlation tests at cosmic distances](#), by Thomas Durt and Jean Schneider, arXiv, April 2026 (10 pages).
- [458] [Proof-of-Concept Experiments for Quantum Physics in Space](#), by Rainer Kaltenbaek, Markus Aspelmeyer, Thomas Jennewein, Caslav Brukner, Martin Pfennigbauer, Walter R. Leeb, and Anton Zeilinger, arXiv, August 2003 (15 pages).
- [459] [Bell's inequality holds for all non-product states](#), by Nicolas Gisin, Physics Letters A, 1991.
- [460] [The classical-quantum boundary for correlations: Discord and related measures](#), by Kavan Modi, Aharon Brodutch, Hugo Cable, Tomasz Paterek, and Vlatko Vedral, Reviews of Modern Physics, 2012.
- [461] [Measures and applications of quantum correlations](#), by Gerardo Adesso, Thomas R. Bromley, and Marco Cianciaruso, Journal of Physics A: Mathematical and Theoretical, 2016.
- [462] [Quantum Discord: A Measure of the Quantumness of Correlations](#), by Harold Ollivier and Wojciech H. Zurek, Physical Review Letters, 2001 (preprint on arXiv).
- [463] [Almost all quantum states have nonclassical correlations](#), by Alessandro Ferraro, Leandro Aolita, Daniel Cavalcanti, Fernando M. Cucchietti, and Antonio Acín, Physical Review A, 2010.
- [464] [Quantum entanglement](#), by Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki, Reviews of Modern Physics, 2009.
- [465] [Steering, Entanglement, Nonlocality, and the Einstein-Podolsky-Rosen Paradox](#), by Howard M. Wiseman, Steve J. Jones, and Andrew C. Doherty, Physical Review Letters, 2007.
- [466] [Bell nonlocality](#), by Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner, Reviews of Modern Physics, 2014.
- [467] [Quantum observables in particle physics](#), by Fabio Maltoni, INFIERI 2025, September 2025 (78 pages).

- [468] [The Uncertainty Principle](#), *Stanford Encyclopedia of Philosophy*, by Jan Hilgevoord, 2006 (14 pages).
- [469] [100 Years of Quantum Uncertainty: From Origins to Modern Insights](#), by Lorcan O. Conlon, Biveen Shajilal, Jie Zhao, Tim C. Ralph, Gerd Leuchs, Ulrik L. Andersen, Syed M. Assad, Ping Koy Lam, and Yunlong Xiao, arXiv, June 2026 (79 pages).
- [470] [Squeezing More from Gravitational-Wave Detectors](#), by Philip Ball, *Physics*, December 2019 .
- [471] [Near-optimal single-photon sources in the solid state](#), by Niccolo Somaschi, Valerian Giesz, Lorenzo De Santis, J. C. Lored, Marcelo P. Almeida, Gaston Hornecker, S. L. Portalupi, Thomas Grange, Carlos Anton, Justin Demory, Carmen Gomez, Isabelle Sagnes, Norberto D. Lanzillotti-Kimura, Aristide Lemaitre, Anthony Auffeves, Andrew G. White, Loic Lanco, and Pascale Senellart, *Nature Photonics*, June 2016 (17 pages)  (preprint on [arXiv](#)).
- [472] [A Quantum Engineer's Guide to Superconducting Qubits](#), by Philip Krantz, William D. Oliver, et al., *Applied Physics Review*, June 2019 (67 pages).
- [473] [Quantum tomography explains quantum mechanics](#), by Arnold Neumaier, arXiv, May 2024 (112 pages).
- [474] [Ideal Projective Measurements Have Infinite Resource Costs](#), by Yelena Guryanova, Nicolai Friis, and Marcus Huber, arXiv, December 2019 (14 pages).
- [475] [Quantum speed limits: from Heisenberg's uncertainty principle to optimal quantum control](#), by Sebastian Deffner and Steve Campbell, *Journal of Physics A: Mathematical and Theoretical*, October 2017 (50 pages).
- [476] [Implementation of multiparticle quantum speed limits on observables](#), by Rui-Heng Miao, Zhao-Di Liu, Chen-Xi Ning, Yu-Cong Hu, Hao Zhang, Chuan-Feng Li, and Guang-Can Guo, arXiv, October 2025 (11 pages).
- [477] [Coherent Quantum Speed Limits](#), by Xuhui Xiao, Hai Wang, and Xingze Qiu, arXiv, March 2026 (5 pages).
- [478] [Quantum speed limits: from Heisenberg's uncertainty principle to optimal quantum control](#), by Sebastian Deffner and Steve Campbell, *Journal of Physics A: Mathematical and Theoretical*, October 2017.
- [479] [How fast can a quantum gate be? Exact speed limits from geometry](#), by Hunter Nelson and Edwin Barnes, arXiv, April 2026 (18 pages).
- [480] [Dynamics of ideal quantum measurement of a spin 1 with a Curie-Weiss magnet](#), by Theodorus Maria Nieuwenhuizen, arXiv, March 2026 (23 pages).
- [481] [Quantum speed limit for measurement probabilities](#), by Agung Budiyo and Sebastian Deffner, arXiv, April 2026 (9 pages).
- [482] [Ultrafast quantum gate operations in a Kramers-Henneberger atom Qubit](#), by A. Tasnim Aynul, C. Li, L. Cruz Rodriguez, and C. Figueira de Morisson Faria, arXiv, August 2026 (30 pages).
- [483] [The uncertainty relation between energy and time in non-relativistic quantum mechanics](#), by L. Mandelstam and I. Tamm, *Journal of Physics (USSR)*, 1945 .
- [484] [The maximum speed of dynamical evolution](#), by Norman Margolus and Lev B. Levitin, *Physica D: Nonlinear Phenomena*, September 1998.
- [485] [Fundamental limit on the rate of quantum dynamics: The unified bound is tight](#), by Lev B. Levitin and Tommaso Toffoli, *Physical Review Letters*, October 2009  (preprint on [arXiv](#)).
- [486] [Bounding the minimum time of a quantum measurement](#), by Nathan Shettell, Federico Centrone, and Luis Pedro García-Pintos, *Quantum*, November 2023.
- [487] [Teaching ideal quantum measurement, from dynamics to interpretation](#), by Armen E. Allahverdyan, Roger Balian, and Theo M. Nieuwenhuizen, arXiv, May 2024 (32 pages).
- [488] [The Quantum Measurement Problem: A Review of Recent Trends](#), by Anderson A. Tomaz, Rafael S. Mattos, and Mario Barbatti, *Philosophical Magazine*, December 2025 (38 pages)  (preprint on [arXiv](#)).
- [489] [A note on the quantum measurement problem](#), by N. David Mermin, arXiv, June 2022 (4 pages).
- [490] [Everything is Entangled in Quantum Mechanics: Are the Orthodox Measures Physically Meaningful?](#), by Christian de Ronde, Raimundo Fernandez Moujan, and Cesar Massri, arXiv, May 2024 (22 pages).
- [491] [Introduction to Koopman-von Neumann Mechanics](#), by Daniel Piasecki, arXiv, December 2021 (123 pages).
- [492] [Measurement of a quantum system with a classical apparatus using ensembles on configuration space](#), by Marcel Reginatto et al., 2022 (25 pages).
- [493] [Koopman-von Neumann Approach to Quantum Simulation of Nonlinear Classical Dynamics](#), by Ilon Joseph, arXiv, September 2020 (19 pages).
- [494] [Heading towards an Algebraic Heisenberg Cut](#), by Mathias Van Den Bossche and Philippe Grangier, arXiv, May 2025 (15 pages).
- [495] [A single quantum cannot be cloned](#), by W. K. Wootters and W. H. Zurek, *Nature*, October 1982 .
- [496] [The no-cloning theorem](#), by William K. Wootters and Wojciech H. Zurek, *Physics Today*, February 2009 (2 pages).
- [497] [No Practical Quantum Broadcasting: Even Virtually](#), by Yunlong Xiao, Xiangjing Liu, and Zhenhuan Liu, arXiv, March 2025 (6 pages).
- [498] [Brief summaries on symmetries in the history of physics-mathematics: James Clerk Maxwell \(1865–1873\), Emmy Noether \(1915–1918\) and Albert Einstein \(1905–1926\)](#), by Raffaele Pisano, *Journal of Physics: Conference Series*, June 2024.
- [499] [Unification of Finite Symmetries in Simulation of Many-body Systems on Quantum Computers](#), by Victor M. Bastidas, Nathan Fitzpatrick, K. J. Joven, Zane M. Rossi, Shariful Islam, Troy Van Voorhis, Isaac L. Chuang, and Yuan Liu, arXiv, April 2025 (30 pages).
- [500] [A Symmetry-Based Taxonomy of Quantum Algorithms](#), by Sakshi Kumar, Sumit Chilkoti, and Mrittunjoy Guha Majumdar, arXiv, August 2025 (5 pages).
- [501] [Low-depth quantum symmetrization](#), by Zhenning Liu, Andrew M. Childs, and Daniel Gottesman, arXiv, May 2025 (20 pages).
- [502] [State preparation and symmetries](#), by Ivana Miháliková, Joseph Carlson, Duff Neill, and Ionel Stetcu, arXiv, October 2025 (10 pages).
- [503] [Discovering Symmetry Groups with Flow Matching](#), by Yuxuan Chen, Jung Yeon Park, Floor Eijkelboom, Jianke Yang, Jan-Willem van de Meent, Lawson L.S. Wong, and Robin Walters, arXiv, December 2025.
- [504] [Invariant Variation Problems \(Invariante Variationsprobleme\)](#), by Emmy Noether, *Nachrichten von der Gesellschaft der Wissenschaften zu Goettingen*, arXiv translation by M. A. Tavel, March 1918 (22 pages).
- [505] [Restrictions on Transversal Encoded Quantum Gate Sets](#), by Bryan Eastin and Emanuel Knill, *Physical Review Letters*, March 2009 (4 pages)  (preprint on [arXiv](#)).
- [506] [Approximate symmetries and quantum error correction](#), by Zi-Wen Liu and Sisi Zhou, arXiv, December 2023 (19 pages).
- [507] [Continuous Symmetries and Approximate Quantum Error Correction](#), by Philippe Faist, *Physical Review X*, October 2020.
- [508] [Time and event symmetry in quantum mechanics](#), by Michael Ridley and Emily Adlam, arXiv, December 2023 (30 pages).

This **Quantum physics basics** subsection contains 170 bibliographical references and notes containing at least 3,458 pages.

3.10 Quantum matter

Quantum matter refers to materials or assemblies of few atoms which, for specific conditions, physical observables such as magnetism, electronic state or optical properties are only described by advanced quantum physics.

They are at the crossroads of statistical physics. Iconic quantum materials are superconducting materials in which, below a certain temperature, electrons behave collectively as a sort of fluid dubbed “quantum fluid”. Other known quantum fluids in physics are superfluid helium, Bose-Einstein condensates, polariton condensates and ultra-cold neutral atoms. They all exhibit quantum mechanical effects at a macroscopic collective level. These phenomena are usually reported at very low temperatures, close to -273°C , and sometimes high pressures but some of them start to emerge in less drastic conditions.

The field of quantum matter was initially related to fundamental physics [509]. Nowadays, however, it is becoming a new field of practical quantum technologies with the emergence of startups like in the quantum batteries domain.

3.10.1 Definitions

Given that all standard matter such as metal, semiconductor or insulator relies on a quantum description, starting with electrons quantum numbers and the atomic structure, how are quantum materials and quantum matter being accurately defined? Where is the frontier? Well, there’s no real consensus on this, a bit like how postulates are formalized in quantum physics.

One of the reasons is that quantum materials range from yet untested theoretical concepts to lab-based experiments, up to industry applications like the ones using various forms of graphene [510]. It is an entirely new research field that is still in the making with a lot of fundamental research.

It is also a field that is hard to dig into, even more than many other fields that are covered in that book, like quantum error correction. So, forgive some of the vagueness of this part. I have not really understood *all* the sentences I wrote here!

The simplest definition I found is “*materials where electrons do interesting things*”. Then, I opened quantum matter’s Pandora’s box and found many other definitions.

The US Department of Energy created its own definition in 2016 with “*solids with exotic physical properties that arise from the interactions of their electrons, beginning at atomic and subatomic scales where the extraordinary effects of quantum mechanics cause unique and unexpected behaviors*” [511] with a slightly simpler one “*solids with exotic physical properties, arising from the quantum mechanical properties of their constituent electrons;*

such materials have great scientific and/or technological potential” [512].

A more precise definition was proposed by Philip Ball in 2017 which is based on materials where electrons are operating collectively as quasiparticles and are frequently confined in some 2D geometries like graphene sheets, with derivatives in 3D assemblies of graphene sheets with small angle rotations called **magic angle**, creating the new field of **twistronics** [513–516].

Quantum materials are also grouped as **strongly correlated materials** where magnetism is important and their behavior is “*dictated by quantum mechanical correlations between electrons*”, and **topological materials** where some symmetry of the material lattice provides protected electronic states on the surface or in the bulk of the crystal [517, 518].

And I didn’t try to find any semantic nuance between quantum matter and quantum materials! In another source [519], quantum matter deals with “*novel phases of matter at zero temperature with exotic properties*”. It adds:

“The main ways of characterizing and manipulating quantum matter are with entanglement, symmetry, and topology:

Entanglement *is the quantum property of correlated physical attributes among particles (position, momentum, spin, polarization).*

Symmetry *refers to features of particles and spacetime that are unchanged under some transformation, seen as the property of a system looking the same from different points of view (a face, a cube, or the laws of physics) and its partner, symmetry breaking (in phase transitions).*

Topology *is the property of geometric form being preserved under deformation (bending, stretching, twisting, and crumpling, but not cutting or gluing). Physical systems may have global symmetric and topological properties that remain invariant across system scales*”. Usually, obtaining some topological order requires cooling at very low temperatures like with superconducting materials.

Classical matter phases transitions are traditionally described with Lev Landau’s symmetry breaking model elaborated in 1937. It describes in a simplified way what happens at phase transitions (like gas $\leftrightarrow$ liquid and liquid $\leftrightarrow$ solid) with the evolution of a symmetry-breaking order parameter (OP) named η (eta). It also describes various types of ordering phenomena like ferromagnetic, ferroelectric, ferroelastic or other types of electronic orders like Mott or insulator-metal transitions systems. In most cases, quantum matter is described by a “topological order” that can’t be explained by Landau’s model. Some examples include topological insulators, topological semimetals, fractional quantum Hall states, quantum spin liquids and Fermi liquids. A Mott transition is a particular type of topological phase transition. Mott insulators are materials that are expected to conduct electricity but are insulators, particularly at low tem-




Figure 126: Quantum wires. Source: [520].

peratures, and under certain conditions which can be controlled, leading to so-called Mott transitions.

As stated before, quantum matter is characterized by being based on collective excitations. These excitations are composite entities that are analogous in their behavior to a single particle named quasiparticles.

It can be quasiparticles that are assemblies of several fermions, mostly electrons and holes, like two electrons in Cooper pairs explaining superconductivity, polaritons, excitons and vortex magnetic phenomena like skyrmions, etc. It can also be collective excitations of bosons like phonons in crystal lattices. There are over 30 identified quasiparticles classes including some that are very exotic and less talked about like the bogoliubon (a quasiparticle found in superconductors) and the wrinklon.

Philip Ball proposes a classification of these quasiparticles in seven categories [513]:

- **Cooper pairs** of electrons in classical superconductivity (high-temperature superconductivity with cuprates requires a more complicated explanation). We cover their various use cases with superconducting qubits and sensors.
- **Relativistic Dirac fermions** such as many-electron excitations in Dirac semimetals and in graphene [521]. Graphene has many applications in sensing and electronics. There was even a European Union Graphene Flagship program launched in 2013 with €1B.
- **Weyl fermions** are massless fermions related to Dirac fermions whose existence was predicted by Hermann Weyl in 1929 and discovered in 2015 at Princeton [522]. These fermions are massless, have a high degree of mobility and are quasiparticle excitations in Weyl semimetals. Topological semimetals could be used in low-consumption spintronic and magnetic memory devices and ultrafast photodetectors [523, 524].
- **Laughlin quasiparticles** proposed by Robert Laughlin in 1983 and who received the Nobel in Physics in 1998 for his theoretical explanation of the fractional quantum Hall effect, together with Horst Störmer and Daniel Tsui, who discovered the effect experimentally. They relate to the “fractional quantum Hall effect” (FQHE, discovered in 1982) in a 2D “electron gas” placed in a magnetic field. It involves electron quasiparticles behaving as if they had a fractional charge, such as $1/3$, $2/5$ or $3/7$, 1 being the charge of a single electron. One use case is to create an electron interferometer [525].
- **Majorana fermions** are hypothetical particles proposed in 1937 by Ettore Majorana, which are their own antiparticles. They exhibit some quantum phenomenon in devices like superconducting nanowires. Their existence is still questioned since, when observed, the related phenomenon measurement can easily be confused with ambient noise. Theoretically, Majorana fermions could lead to creating topological qubits quantum computers with a better resistance to quantum noise and errors and some better scalability potential in the fault-tolerant regime (FTQC).
- **Anyons** were introduced in 1977 by Jon Magne Leinaas and Jan Myrheim and named in 1982 by

Frank Wilczek [526]. In two dimensions, the exchange of two identical particles is no longer restricted to the $+1$ of bosons and the -1 of fermions. Abelian anyons acquire an arbitrary phase $e^{i\theta}$ interpolating continuously between the bosonic and the fermionic cases, while non-Abelian anyons have no single exchange phase at all and instead transform within a degenerate state space under braiding, the exchange being described by a matrix rather than by a number. Evidence for fractional Abelian exchange statistics in selected fractional quantum Hall states were found experimentally in 2020 in two-dimensional electron gases at filling factor $1/3$, first with an anyon collider at ENS Paris and C2N, then with a Fabry-Perot interferometer at Purdue University [527, 528]. These results do not cover every anyon class and non-Abelian statistics remain unobserved to date. Anyons could also show up in quantum spin liquids [529], which can appear in magnetic materials where electron spins are not orderly aligned but are entangled. No material is yet an undisputed quantum spin liquid, the best candidates being α - RuCl_3 and the herbertsmithite kagome lattices, where 2020 measurements are compatible with, but do not prove, a spin liquid state [530]. It could help to create innovative electronic memories. This state of matter was envisioned in 1973 by Philip W. Anderson with his resonating valence bond model [531, 532].

- **Skyrmions** take the form of vortex-like topological quasiparticle excitations of spins in some magnetic materials. They were envisioned in 1962.

We could still add here various classes and subclasses of quantum materials:

- **Spin glasses** where electron spins freeze in a disordered fashion at some non-zero temperature. It leads to the notion of **quantum glasses** [533].
- **Plasmons** which are the quanta of collective oscillations of an electron density. They exist in the bulk of a conductor, at a bulk plasma frequency in the few eV range for ordinary metals, and also at interfaces as surface plasmons. When a surface plasmon hybridizes with an electromagnetic wave, the resulting coupled charge-density and field mode is a surface plasmon polariton, which is the object actually exploited in plasmonics. It could also help create energy savings and faster data storage solutions.
- **Topological insulators** are materials whose bulk is insulating and whose boundary, edges in 2D and surfaces in 3D, hosts symmetry-protected conducting states [534, 535]. In a quantum spin Hall edge, opposite spin channels counterpropagate, spin and momentum being locked together. These helical channels carry charge current. What the topology brings is that spin-momentum locking suppresses elastic backscattering by nonmagnetic disorder. Topological insulators could for example lead to the creation

of new breeds of energy-saving and fast-switching transistors [536].

- **Quantum wires** are conducting wires with quantum confinement effects modifying the transport properties, mostly when the wires have a diameter of a few nanometers, even down to a single atom. They are usually called nanowires. Carbon nanotubes are a class of quantum wires (shown in Figure 126).
- **Spin-torque materials** that are already used in some low-power non-volatile magnetic memories (STT-RAM or STT-MRAM).
- **Time crystals** which we'll cover later, and it is the source of a lot of headaches.
- **Wigner crystals** are another very weird phenomenon. Predicted by Eugene Wigner in 1934 (the same Wigner of the Wigner function used in quantum photonics), it consists of crystals made of electrons, of course also at very low temperatures. They were experimentally observed in 2018 by an Israeli-US-Hungarian team in one dimension at 10 mK using carbon nanotubes for their measurement [537] and 2020 in 2D at 80 mK by a team from ETH Zurich [538] (Figure 129).
- **Quantum batteries** are still at the proof-of-concept stage, with laboratory prototypes storing microjoules for microseconds. Collective effects would in principle give them a shorter recharging cycle than traditional batteries.




Figure 127: Quantum dots used in LCD screen and lighting. In TV LCD screens, these dots are used to convert the blue light coming from the LED into pure green and red light, which increases the “gamut” of the display, enabling it to display a richer set of colors thanks to better saturated primary colors. Source: Olivier Ezratty, 2023-2025, inspired from: [539], with a correction on the emission wavelengths per quantum dot size.

- **Quantum dots** that are used in LCD screens are semiconductor nanocrystals whose size-dependent emission comes from the quantum confinement of electron-hole pairs, an effect that is entirely quantum. They are nonetheless usually classified as nanomaterials rather than within the strongly correlated and topological “quantum materials” class used in this part, which is a taxonomy convention and not a claim that their mechanism is classical. They are made of powder with tiny compound grains of different sizes between 2 and 6 nm which are used to down-convert the blue light coming from LEDs into red and green light, creating a better balanced coverage of primary colors, as shown in Figure 127. They were first observed in 1981 by Alexei Ekimov in doped glasses in Russia, explained in 1982 by Alexander Efros, also from Russia, and obtained in colloidal solution by Louis Brus in 1983 [540, 541]. Ekimov, Brus and Moungi Bawendi shared the 2023 Nobel prize in chemistry for this discovery and for the controlled synthesis of quantum dots. The main problem is to replace cadmium that is a pollutant.

These LCD screens quantum dots must not be confused with the quantum dots used in silicon qubits to trap single electrons and control their spin as well as the quantum dots used in unique photon sources like the ones from Quandela.

Some other concepts related to quantum matter mandate some explanations:

Many quantum matter species happen in crystals. And there are a lot of types of crystals classified by their crystallographic order! There are 230 crystallographic space groups organized in 7 crystal systems named triclinic, monoclinic, orthorhombic, tetragonal, trigonal, hexagonal and cubic and subclasses with primitive centering, centered on a single face, body centered and face centered (Figure 128) [542].




Figure 128: Bravais lattices and crystal structure classification. Source: Wikipedia.

One key notion in crystallography is chirality which describes how crystal structures break spatial symmetry and are not identical to their mirrored structure [543, 544]. There are also 1,651 magnetic space groups which

describe magnetism configurations at the atom level in crystal lattices [545–547] (Figure 130).

Another key notion in quantum matter is time reversal symmetry. A time reversal symmetry means that the material looks the same when looking at a time scale backwards and forward.

time reversal symmetry

do not change with time reversal	changes with time reversal
position of particle in space	time of events
particle acceleration in space	particle velocity
force on the particle	particle linear momentum
particle energy	magnetic field
electric potential and field	electric current density
density of electric charge	power / rate of work done
energy density of the EM field	

Figure 131: Time reversal symmetry explained. It doesn't mean some quantum object goes back in time and relates more to time symmetries. Under the conventional antiunitary time reversal operator, the electric field $\mathbf{E}$ is even and the magnetic field $\mathbf{B}$ is odd, while momentum $\mathbf{p}$ and current density $\mathbf{J}$ are odd.

Reversing time means looking backwards in time only from a mathematical standpoint, not physically reversing time. There's no way you can change the arrow of time backwards. Time reversal is not a time machine! Figure 131 presents an inventory of some physical properties that change or do not change with time reversal. Some experiments of time reversal are also implemented with photons, but the related symmetries do not mean these systems can really go back in time like in science fiction [548, 549] (Figure 131).

3.10.2 Superconductivity

Superconductivity occurs when, below a critical temperature and also below critical magnetic field and critical current thresholds, some conducting materials carry a DC electric current with zero electrical resistivity. In an ordinary conductor, charge carriers lose momentum by scattering off lattice vibrations, meaning phonons, and off impurities, defects, grain boundaries and other electrons, transferring energy to the lattice as heat. This is the Joule effect.




Figure 132: Some materials have zero resistivity below a threshold temperature (T_c). © Olivier Ezratty, 2020-2025.




Figure 129: Wigner crystals. Source: [538].




Figure 130: Ferromagnetism, paramagnetism, ferrimagnetism and antiferromagnetism explained. Source: Wikipedia.

Discovery. Superconductivity was discovered experimentally in 1911 by **Heike Kamerlingh Onnes** (1853-1926), Cornelis Dorsman, Gerrit Jan Flim and Gilles Holst at the University of Leiden in the Netherlands, with solid mercury at 4.2K. Kamerlingh Onnes also discovered that a magnetic field whose intensity depends on temperature could make the superconducting effect disappear [550]. The critical temperature T_C of a superconductor is the maximum temperature at which it shows superconductivity behavior.

Explanation. Superconductivity interpretation was formulated much later, in 1957 and achieved by **John Bardeen**, **Leon Neil Cooper** and **John Robert Schrieffer** of the University of Illinois. They built the so-called **BCS theory** which explains how superconductivity works [551–553]. Before the BCS theory, many physicists had tried unsuccessfully to explain superconductivity: Albert Einstein, Niels Bohr, Lev Landau, Max Born, Felix Bloch, Léon Brillouin, John Bardeen (co-inventor of the transistor), Werner Heisenberg and Richard Feynman. John Bardeen was awarded twice the Nobel prize in physics, one in 1956 for the invention of the transistor with William Shockley and Walter Brattain and the other for the interpretation of superconductivity in 1972 with Leon Neil Cooper and John Robert Schrieffer. Cooper co-created the BCS theory

at the age of 27 and won the corresponding Nobel Prize at the age of 42.




Figure 133: Superconductivity can be explained by the creation of Cooper pairs of electrons with opposite spins flowing in crystal lattice, creating bosons © Olivier Ezratty, 2020-2025.

In the BCS theory, the exchange of virtual phonons creates an effective attractive interaction between electrons of opposite momentum and opposite spin near the Fermi surface, which overcomes their Coulomb repulsion below T_C (Figure 133). This binds them into Cooper pairs. Conventional s-wave pairs are spin singlets with a total spin of zero, which is what lets the condensate behave as a single macroscopic quantum state, but this does not generalize, unconventional superconductors having triplet or other pairing structures. Cooper pairs can

also be formed with atoms as with helium 3, a fermion, in its superfluid state named a fermionic condensate.

The **Ginzburg-Landau** theory was published jointly by Vitaly Ginzburg and Lev Landau in the USSR in 1950, seven years before BCS. It is a phenomenological theory providing a macroscopic description of superconductivity, focusing on the collective behavior of superconducting electrons rather than on the microscopic details of Cooper pair formation as in BCS. It focuses on the order parameter which describes the superconducting state and the transition from the normal state to the superconducting state. It helps describe the Meissner effect, the formation of vortices and to understand the behavior of superconductors near the critical temperature and in the presence of external magnetic fields. It also predicts the two types of superconductors (type I and II) that we will describe later. Lev Gor'kov showed in 1959 that Ginzburg-Landau is a limiting form of BCS near T_C , which turned a phenomenological model into a controlled approximation [554]. Ginzburg was awarded the Nobel prize in physics in 2003 for this work, together with Alexei Abrikosov and Anthony Leggett.

The electrons in Cooper pairs are not necessarily close together since interaction is long range. Paired electrons may be between 100 nm and 1 μm apart, a distance that is usually greater than the average interelectron distance so that many Cooper pairs can occupy the same space. This distance is the coherence length of the Cooper pair. It is large compared with the mean separation between conduction electrons in a metal that is about 0.15 nm for copper. As a result, Cooper pairs overlap greatly.

Elements. About 30 chemical elements are superconducting at ambient pressure and about 20 more only under high pressure, with very variable temperature and pressure thresholds (Figure 138). T_C does not grow monotonically with the number of valence electrons. Bernd Matthias's empirical rules place the T_C maxima of transition metals at roughly 4.7 and 6.5 valence electrons per atom. In general, metals that are superconductors are poor conductors in their normal state and most good conductors like copper, gold and silver are not superconductors, since the strong electron-phonon coupling which degrades normal conduction is precisely what favors pairing. Superconductivity is also possible with composite alloys such as niobium-germanium (Nb_3Ge , 23K) and niobium-titanium, or with copper-based ceramics (cuprates). The most common superconducting materials are aluminum and a niobium and titanium alloy, used in superconducting wires in MRI systems and superconducting qubit cryostats. The superconducting properties of the niobium-titanium alloy were discovered in 1962. It is widely used in the cooling of MRI scanners but also in many scientific instruments, notably in the ITER experimental nuclear fusion reactor at Cadarache in France. Titanium becomes superconducting at 390 mK, aluminum at 1.2K, indium at 3.4K and niobium at 9.26K.

Later experiments and extrapolations on the persistence of circulating currents injected into macroscopic superconducting rings found that the lower bound of these permanent currents was around 10^5 years.

Meissner effect. Superconductivity explains unexpected phenomena such as the levitation of magnets above superconductors immersed in liquid nitrogen. Superconducting ceramics, discovered in 1986, can be used in this striking experiment.




Figure 134: The Meissner effect.

The magnetic field is then expelled from inside the superconducting material. This is the Meissner effect, discovered in 1933 by **Walther Meissner** (1882-1974, German) and his collaborator **Robert Ochsenfeld** (1901-1993, German), which is why it is also named the Meissner-Ochsenfeld effect [555] (Figure 134 and Figure 135).

Type I superconductors are characterized by having a single critical magnetic field, denoted as B_C (C for “critical field”) as shown in Figure 137. Below this critical field, an ideal type I superconductor expels magnetic flux from its bulk and exhibits perfect diamagnetism (being pushed out of regions with stronger magnetic fields). The screening is not absolute. The field decays exponentially over the London penetration depth, typically 20 to 100 nm, within which the screening supercurrents flow, and above B_C superconductivity is destroyed altogether. A finite sample with an unfavorable demagnetization factor also develops an intermediate state with coexisting normal and superconducting domains well below B_C . Type I superconductors are typically s-wave superconductors with isotropic behavior. A complete Meissner effect over the whole superconducting phase only happens with type I superconductors, type II ones expelling the flux entirely only below B_{C1} .

Type II has an intermediate phase between the classical metallic phase and the superconducting phase that allows the magnetic field to pass partially, with two critical magnetic fields: the lower critical field (B_{C1}) and the upper critical field (B_{C2}). Below B_{C1} , type II superconductors behave similarly to type I superconductors, expelling magnetic flux.

Type I and Type II. The first classification of superconducting relates to the way they react to an ambient magnetic field. Type I and II superconductors are mathematically and quantumly explained by the **Ginzburg-Landau** theory created in 1950 [556] (Figure 137).




Figure 135: Meissner effect explanation.

Between B_{C1} and B_{C2} , they allow partial penetration of the magnetic flux into the material while still exhibiting superconducting behavior. Above B_{C2} , the superconducting material transitions into the normal state and loses all its superconducting properties.

There is even a proposal for a so-called Type III superconductor which shows up in thin films with different wave function mechanisms, “described by a topological gauge theory and corresponds to an inhomogeneous network of condensate droplets getting connected by tunneling pairs percolation and is destroyed by vortex liberation instead of pairs breaking” [557, 558]! This is a specialized proposal for a particular thin-film regime and not a generally accepted third class sitting alongside type I and type II, which are separated by the value of the Ginzburg-Landau parameter κ relative to $1/\sqrt{2}$.




Figure 137: Type I and II superconductors characteristics. They respond differently to magnetic fields. Type I are typically pure metals like lead with a complete Meissner effect. They lose superconductivity abruptly when exposed to small magnetic fields, making them suitable only for low-field environments. Type II such as niobium-titanium alloys (NbTi, used in cabling below 4K in superconducting qubit cryostats) or high-temperature ceramic compounds tolerate stronger magnetic fields. It makes Type II materials relevant in MRI scanners, particle accelerators, and maglev trains. Source: [Critical Magnetic Field](#), 1999.




Figure 136: Representation of the Cooper pair wave function for s-wave, p-wave and d-wave superconductors. [Image source](#). 2023.

s-wave, d-wave and p-wave. Superconducting materials are also classified in three kinds, dealing with the symmetry of the wave function describing the Cooper pairs. The letters come from the spectroscopic notation of atomic orbitals (sharp, principal, diffuse) and label the orbital angular momentum ℓ of the pair, respectively 0, 1 and 2. That labeling is exact only for a system with continuous rotational symmetry. A crystal has none, so in a real superconductor these letters designate how the gap function transforms under the crystal point group rather than a well-defined total pair angular momentum. It has implications on the anisotropy of various properties and to the materials applications.

In s-wave superconductors, the electron pairs have an angular momentum of zero. Their wave function is invariant under any rotation, thus the spherical representation shown in Figure 136 and them being isotropic, i.e., their properties being the same in all directions. It allows the electrons to move freely without scattering and leads to the observed zero resistance and expulsion of magnetic fields shown in the Meissner effect.

d-wave superconductors have a gap that transforms with d-wave symmetry under the crystal point group. For the common $d_{x^2-y^2}$ form, the gap function varies as $\cos(2\varphi)$, so it changes sign under a 90-degree rotation and has four lobes separated by nodes where the superconducting gap vanishes. Those nodes leave low-energy excitations at any temperature, which shows up as power-law rather than exponential temperature dependences in the specific heat and the penetration depth. Their behavior can be anisotropic with properties varying depending on the direction in the material. Certain high-temperature superconducting materials such as the cuprates are usually of the d-wave type.

p-wave superconductors have a gap of odd parity under inversion, which changes sign under a 180-degree rotation and, by the Pauli principle, goes together with spin-triplet pairing. They also exhibit anisotropic be-

superconducting elements																	
1 H hydrogen	2 He helium															18 He helium	4.003
3 Li lithium	4 Be beryllium															13 B boron	10.81*
11 Na sodium	12 Mg magnesium	13 Al aluminium	14 Si silicon	15 P phosphorus	16 S sulfur	17 Cl chlorine	18 Ar argon										
19 K potassium	20 Ca calcium	21 Sc scandium	22 Ti titanium	23 V vanadium	24 Cr chromium	25 Mn manganese	26 Fe iron	27 Co cobalt	28 Ni nickel	29 Cu copper	30 Zn zinc	31 Ga gallium	32 Ge germanium	33 As arsenic	34 Se selenium	35 Br bromine	36 Kr krypton
37 Rb rubidium	38 Sr strontium	39 Y yttrium	40 Zr zirconium	41 Nb niobium	42 Mo molybdenum	43 Tc technetium	44 Ru ruthenium	45 Rh rhodium	46 Pd palladium	47 Ag silver	48 Cd cadmium	49 In indium	50 Sn tin	51 Sb antimony	52 Te tellurium	53 I iodine	54 Xe xenon
55 Cs caesium	56 Ba barium	57-71 La-Lu	72 Hf hafnium	73 Ta tantalum	74 W tungsten	75 Re rhenium	76 Os osmium	77 Ir iridium	78 Pt platinum	79 Au gold	80 Hg mercury	81 Tl thallium	82 Pb lead	83 Bi bismuth	84 Po polonium	85 At astatine	86 Rn radon
87 Fr francium	88 Ra radium	89-103 Ac-Lr	104 Rf rutherfordium	105 Db dubnium	106 Sg seaborgium	107 Bh bohrium	108 Hs hassium	109 Mt meitnerium	110 Ds darmstadtium	111 Rg roentgenium	112 Cn copernicium	113 Nh nihonium	114 Fl flerovium	115 Mc moscovium	116 Lv livermorium	117 Ts tennessine	118 Og oganesson
57 La lanthanum	58 Ce cerium	59 Pr praseodymium	60 Nd neodymium	61 Pm promethium	62 Sm samarium	63 Eu europium	64 Gd gadolinium	65 Tb terbium	66 Dy dysprosium	67 Ho holmium	68 Er erbium	69 Tm thulium	70 Yb ytterbium	71 Lu lutetium			
89 Ac actinium	90 Th thorium	91 Pa protactinium	92 U uranium	93 Np neptunium	94 Pu plutonium	95 Am americium	96 Cm curium	97 Bk berkelium	98 Cf californium	99 Es einsteinium	100 Fm fermium	101 Md mendelevium	102 No nobelium	103 Lr lawrencium			

Figure 138: Table of chemical elements with those which are superconducting. Source: Wikipedia and various other sources.

havior, similar to d-wave superconductors, with different anisotropy patterns.

There are also a lot of so-called “unidentified superconducting objects” with weird characteristics like palladium hydrides, multi-walled carbon nanotubes (6.8K), moiré graphene sheets (twistronics) and doped graphite. Measured PdH sits around 9K and the 52K to 133K figures come from first-principles predictions on cubic MH₄ hydrides at ambient pressure, which are not experimentally confirmed to date [559]. Over a thousand candidate superconductor materials have been inventoried [560, 561].

Use cases. Here are the several known industry use cases for superconductors.

MRI scanners (Magnetic Resonance Imaging, derived from nuclear magnetic resonance, the “nuclear” being dropped in the clinical naming) use large superconducting magnets that are cooled with liquid helium. Scanners are encased in a protective coating to constrain the magnetic field inside the scanner. The niobium-titanium coil wiring is enveloped in a copper matrix (Figure 139).

CERN LHC is also using this combination in large physics instruments with 1,200 tons of cables including 470 tons of NbTi (niobium-titanium), the rest being copper, for about 7,600 km of installed cable around the 27 km ring. During nominal Run 2 and Run 3 operation, superconductivity carries a current of 11,850 A

generating a powerful magnetic field of 8.33 tesla creating a centripetal force holding the accelerated particles. These magnets are cooled at 1.9K by an inventory of about 120 tons of helium, the superfluid helium-4 bath surrounding a cold mass of 36,000 tons. The initial cooldown alone vaporizes about 10,000 tons of liquid nitrogen in the heat exchangers that first bring the helium down to 80K. Their cables are made of niobium-titanium filaments surrounded by copper. The LHC cryogenic plant alone draws about 40 MW. The whole CERN site, accelerators, detectors and infrastructure included, consumed 1,215 GWh in 2022 and sits around 1.2 to 1.3 TWh per year in beam years, the LHC accounting for roughly 55% of it, so about 650 to 700 GWh, according to CERN’s environment reports. It is the largest and most powerful refrigerator in the world.

Transportation. The Chuo Shinkansen Maglev high-speed train in Japan, which has been undergoing trials since 2013 and is expected to reach a commercial speed of 505 km/h uses a superconductive based magnetic suspension with a rather expensive infrastructure. Power consumption per passenger/kilometer is three times that of traditional Shinkansen, but it is still competitive with airplanes. The 286 km Tokyo-Nagoya line was initially planned for commercial service in 2027, but JR Central abandoned that date in March 2024 since tunnel drilling in Shizuoka prefecture had not started. The opening is now expected in 2034 at the earliest, with JR Central



Figure 139: MRI in hospitals have been common place for decades. They used a strong magnet cooled by liquid helium.

mentioning 2035, and the Osaka extension has no firm timeline.

Electric motors and generators. Superconductivity has also been studied to improve the efficiency of electric motors and generators with HTS Synchronous Motors (High-Temperature Superconductors) [563]. It allows a reduction of motors size and efficiency improvements. It is based on superconducting materials that only require liquid nitrogen cooling, but some systems still use helium-based cooling. Studies began in the 1980s and these engines and generators are beginning to be deployed in the military navy and in wind power generation, notably at **ASMC**, **Sumitomo Electric** [564] and with the European **EcoSwing** project, which involves Sumitomo's cryostat division.

Electric cables. HTS superconducting cables can transport electricity without power loss and greater capacity to meet the ever-increasing demand. Several pilot projects have been launched in the USA, Germany and South Korea [565]. One 600 m long underground transport cable was deployed in Long Island in 2008 with the participation of Nexans (France). It can supply electricity to 300,000 homes [566]. HTS cables current transport relies on a thin 4 μm layer (Figure 141). The project cost was \$46.9M.

In South Korea, KEPCO (Korea Electric Power Corp) deployed in 2019 a 23 kV HTS cable between the 154 kV Shingal and Heungdeok substations. The Chicago Resilient Electric Grid (REG) went live in 2022 with a 200 m test HTS cable, with the goal of deploying a 5 km loop connecting downtown Chicago substations in a second stage. The SWM SuperLink cable project in Munich, Germany, created a HTS test infrastructure between 2021 and 2024 for €5M, with a potential deployment of a 12 km transport line in the city [567, 568], the HTS cable being provided by NKT Cables Group (Denmark).




Figure 140: MRI principle. Source: [562].




Figure 141: Nexans superconducting cable.

Superconducting qubits exploit the Josephson junction which is usually a sandwich with aluminum oxide surrounded by aluminum. This technology is used in superconducting qubits, in quantum sensing with SQUIDs (superconducting quantum interference device) and in superconducting electronics (see an excellent review in [569]). Josephson junctions have a relationship between voltage and frequency which enables the creation of various sensors. It can convert a voltage to frequency as well as a frequency to voltage (with the inverse AC Josephson effect using a microwave impulse). We also find it in the type II niobium-titanium based superconducting cables used for reading the state of superconducting and electron spin qubits. Beyond quantum computing, superconducting qubits are useful in multiple domains in condensed matter physics and even in the search of dark matter [570].

We may wonder why scientists are not using high-temperature superconductors to build superconducting qubits? There are various reasons, one being the need to avoid thermal photons that can destroy the qubit coherence. Also, superconductors used in a Josephson junction must be of the s-wave type due to their ro-

bustness against magnetic fields, well controlled energy levels and transition frequencies used in qubit control, their ease of manufacturing on thin films, and electrons isotropic pairing symmetry which simplifies the design of the qubits. Still, some low-TRL research is going on for the design of cuprates based superconducting qubits, aka “flowermon” [571] or for the creation of SQUIDs used in quantum sensing [572].

Superconducting electronics. Superconductivity could also be used to create processors operating at low temperatures. Josephson junction circuits switch extremely fast, a rapid single flux quantum (RSFQ) T-flip flop having been demonstrated at up to 770 GHz at 4.2K as early as 1999 [573, 574]. A switching rate is however not a cross-architecture performance metric and cannot be placed side by side with the 4 to 5 GHz peak clock of a regular server CPU. Useful throughput depends on logic density, on-chip and off-chip memory, I/O bandwidth and cryogenic overhead, the latter costing on the order of 10^3 W of room-temperature power per watt dissipated at the 4K stage. None of these have been demonstrated at CMOS scale with RSFQ. An MIT team announced in July 2019 a proposal for a technique to create spiking neurons with superconducting Josephson effect circuits using nanowires [575]. This is still a research field with very few industry applications at this point. We’ll look at this field in a specific section on unconventional computing. Superconducting electronics could be very useful to control superconducting and electron spin qubits.

High-Temperature Superconductors (HTS) also named HTC (for “high T_c ” or high-critical temperature superconductors). **Georg Bednorz** and **Alex Müller** from IBM Zurich opened the field in 1986 with a lanthanum barium copper oxide ceramic superconducting at 35K, which earned them the Nobel prize in physics the following year [576]. The 77K barrier (-196°C , the temperature of liquid nitrogen) was crossed in 1987 by Maw-Kuen Wu and Paul Chu with $\text{YBa}_2\text{Cu}_3\text{O}_7$ (YBCO) at 93K [577]. These materials are ceramic oxides and not metal alloys. There are about three sorts of HTS: rare-earth barium copper oxide (ReBCO) aka cuprate ceramics, magnesium diboride MgB_2 with T_c of 39K which was discovered in 2001 [578] and iron based superconductor materials (IBS), discovered in 2008 by Hideo Hosono’s team at 26K and raised to 55K within months [579–581]. Only the cuprates operate above liquid nitrogen temperature, MgB_2 and the IBS being “high temperature” only relative to the 4K of NbTi.

A record was achieved in 2019 with a molecule combining lanthanum and hydrogen (LaH_{10} , illustrated in Figure 142) at 250K, thus -23°C and a near-ambient temperature. In the latter case, however, it works at a huge pressure of about 170 GPa, representing more than 1.7 million times the atmospheric pressure, which is 101,325 Pa [582, 583]. A parallel measurement reported an onset above 260K between 180 and 200 GPa [584]. Another record was broken in 2019 with YH_6 (yttrium

hydride) at 224K and 166 GPa [585]. Then, in 2020, a CEA team obtained metallic hydrogen at 425 GPa, but that experiment evidenced an insulator-to-metal optical transition and not superconductivity [586, 587]. A record of 15°C at 267 GPa was announced in the USA in 2020 with a carbonaceous sulfur hydride [588], but this Nature paper by Ranga Dias was retracted in 2022 and the record no longer stands. A less impressive 2022 record was created in China with clathrate calcium hydride (CaH_6) being superconducting at 215K and 172 GPa [589].

You always see this trade-off between superconducting temperature and pressure. At this very high pressure, practical use cases are not easy to implement! But at lower temperatures, interesting used cases arise like with single photons detectors [590]. Hence the willingness to use quantum simulators or computers to run superconductivity quantum equations and identify materials that would be superconducting at room or near-room temperature. Another fancy solution is about lowering the room temperature [591].




Figure 142: LaH_{10} is a high superconducting temperature molecule, but at a very high pressure of about 170 GPa.

We then have the quest for ambient pressure and temperature superconductors. The categories of evidence at play are worth separating before reading any of these claims, since they carry very different weight. A first-principles prediction of a superconducting phase, a measured resistance drop alone, a measured diamagnetic or Meissner signature, an independent replication and a journal retraction are five distinct things, and most publicized room-temperature claims never went past the first two. Some were announced in the past. Like with Indian Institute of Science (IISc)’s questionable silver nanoparticles in a gold grid (not cheap) [592, 593] or Ranga Dias’s papers that were retracted in 2022 (Nature, carbonaceous sulfur hydride), in August 2023 (Physical Review Letters, manganese sulfide) and in November 2023 (Nature, lutetium hydride). A University of Rochester investigation concluded in 2024 that the experimental data had been fabricated [594, 595].

LK99 which appeared in July 2023 made the news as being a potential ambient room and pressure supercon-

ducting material. The copper-lead-phosphorus-oxygen material was discovered by a team of South Korean researchers in July 2023 [596, 597]. It drew a strong interest since, if working as advertised, it would be the Holy Grail of superconductivity and have significant industry applications as showcased above in this text. Patents were filed and a startup even already created. How do you determine that a material is superconducting? Many tests must be done on AC susceptibility, temperature-dependent critical field and critical current, single-particle tunnelling gap, jump in specific heat at T_C , Josephson tunnelling, AC Josephson effect, Meissner effect and so on (Twitter source). The original papers were quite sketchy in the first place. The claimed superconducting effect was limited to very weak currents and voltages (under 1V and under 500 mA). The authors also described LK99 as being of d-wave type, a claim that was never independently established.

Very quickly, other laboratories in China and the USA, tried to reproduce the experiment [598]. It did generate very mixed to bad results. A scientist from the Lawrence Berkeley National Laboratory used a classical DFT-based simulation to explain the potential origins of the material superconductivity [599]. One China team proposed a way to improve the performance of the material [600]. Another China team couldn't replicate room temperature superconductivity at ambient temperature [601] and another one could make it but only at 100K, which is still kind of cold [602]. Another experiment did show the effect of ferromagnetism explaining the observed levitation, but not superconductivity. But the material was produced with sintering instead of annealing which may explain the different results [603]. Other papers explained that the material was only diamagnetic, which is commonplace [604] and that the near zero resistivity at 387K was explainable by some Cu_2S impurities [605]. At last, the Korean academy temporarily invalidated the work [606]. You can also find a lot of more or less valid scientific information about LK99 in various community sites [607–609] and even an estimation of the quantum computing resources that would be needed to simulate LK99 [610].

Just after LK99 was announced, a weird US company, Taj Quantum, specialized in securing communications with a Blockchain, announced the validation of a patent for a type II ambient temperature superconducting material using graphene doped with aliphatic hydrocarbons, created by John Wood and Paul Lilly. The conclusion in August 2023 was that LK-99 is “an insulator with a resistance in the millions of ohms, too high to run a standard conductivity test. It shows minor ferromagnetism and diamagnetism, but not enough for even partial levitation” [611]. Case closed!

Another 2022 paper published in a lesser known peer-reviewed journal in December 2023 with the participation of Terra Quantum and a University in Brazil also claimed to have discovered a graphene-based room temperature superconductor [612]. It got less attention

due to “room temperature superconductor” fatigue. It still drove some skeptical views from other specialists [613].

This showcases the problem with today's scientific communication with the media and other influencers broadcasting news before scientists can carefully examine new claims. The timing discrepancy between these two categories of players is a given, teaching us to always be careful about extraordinary claims. But when the scientific community is mobilized to fact check claims, it can make miracles in a very short time.

3.10.3 Superfluidity

Superfluidity is yet another quantum physics phenomenon to cover here. Its best known incarnation is superfluid helium which, at ambient pressure, never freezes, no matter how low the temperature can be. It also shows up in ultracold atomic Bose and Fermi gases, in exciton-polariton condensates and, according to models, in the inner crust of neutron stars.




Figure 143: Helium superfluidity phase diagram. Source: Wikipedia.

Superfluid liquid has zero viscosity and flows without any loss of kinetic energy. When poured into a recipient, it tends to rise by capillary action on its rim and flow out of it. It can even pass through very fine capillaries (Figure 144).

Helium was first liquified in 1908 at 4.2K by Heike Kamerlingh Onnes, the discoverer of superconductivity in 1911. Its superfluidity was highlighted independently in 1938 by **Pyotr Kapitsa** (1894-1984, USSR, who was awarded the Nobel Prize in 1978 for his work in the field of low temperatures), **John Frank Allen** (1908-2001, USA) and **Don Misener** (1911-1996, USA) [614].

There are two isotopes of helium: ^3He with a single neutron, which is the least abundant in nature, and ^4He , with two neutrons, the most common. The latter is a boson, with an integer spin, giving it different properties from helium 3, which is a fermion with a half-integer spin. At low temperature, ^4He is a bosonic but strongly interacting quantum liquid. Its superfluid phase does involve Bose condensation, but only about 8% of the atoms sit in the zero-momentum condensate at $T = 0$,




Figure 144: Superfluidity effect, when the liquid helium is climbing the surface of its recipient to escape it. Source: [615].




Figure 145: Visualization of the superfluidity phenomenon in a curved recipient. Liquid helium also escapes it from the edge. Source: [615].

against close to 100% for a dilute atomic BEC, so the two must not be equated. ^3He behaves differently, being fermions, and assemble in pairs similar to electron Cooper pairs. It becomes superfluid at lower temperatures than ^4He , at around 1 mK in the absence of a magnetic field (see the phase diagram in Figure 143), vs. 2.17K for ^4He .




Figure 146: A dilution refrigerator principle. It is using the two isotopes of helium, ^3He and ^4He which have different phase transitions and are both superfluid at low temperature. A primary circuit cools the circuit at 4K with ^4He , on the top left. The secondart circuit called a dilution, is fed by ^3He gas and has the heavier ^4He at its bottom in the mixing chamber. [Wikimedia](#).

Its superfluidity was only discovered in 1972. David Morris Lee (1931), Douglas Dean Osheroff (1945) and

Robert Coleman Richardson (1937-2013) were awarded the Nobel prize in physics in 1996 for their discovery of helium-3 superfluidity. The different properties of ^3He and ^4He are used to operate the dilution cryogenics systems that equip many quantum computers whose operating temperature is between 10mK and 1K.

Superfluid helium could also be used to create low dissipation mechanical resonator-based dark matter detectors [616]. It is also used for the creation of topological effects [617].

Industrial demand for helium is spread across many industries: medical imaging for MRI systems magnets cooling, then microelectronics industries (Figure 147).




Figure 147: Usages of helium, the main one being MRI in hospitals (20%). Source: [Edison Investment Research](#), February 2019.

3.10.4 Bose-Einstein Condensates

Bose-Einstein condensates is a particular state of matter. These are extremely low-density gases of bosons cooled down to very low temperatures, at the lowest energy level we can set matter in, below solid state. Rubidium-87 and sodium-23 are the most famous elements that were experimented in this matter state. Superfluid ^4He is a related but different case, being a dense and strongly

interacting liquid where only about 10% of the atoms sit in the condensate.

It took a while between the work of Bose and Einstein in 1924 and the experimental discovery of BECs in 1995 by **Carl Wieman** and **Eric Cornell** at JILA with rubidium 87 at 170 nK, followed four months later by **Wolfgang Ketterle** at MIT with sodium 23. It was cooled with laser-based Doppler effect and magnetic evaporating technique. The three shared the 2001 Nobel prize in physics.

BECs play an important role in quantum technologies. They led to the control of individual atoms that are used in quantum simulators and in quantum gravimeters, although at a higher temperature and with atoms being exploited as individual quantum objects in the case of computing. Together with superfluids and supersolids, BECs belong to the field of quantum hydrodynamics.

BECs are even experimented in space. The NASA Cold Atom Laboratory (CAL) has been studying ultra-cold gases in microgravity conditions, producing condensates at about 100 nK and reaching effective expansion temperatures in the picokelvin range, with a below 100 pK design goal, including Bose-Einstein condensates with rubidium-87 and potassium-41 in the ISS since May 2018 [618].

Very low temperature can be the playing field of various phenomenon, including the association of atoms in many-body objects like trimers (three atoms) as described by Efimov physics [619].




Figure 148: Bose-Einstein condensates visualized at 400, 200 and 50 nK. Source: [620].

3.10.5 Supersolids

Supersolidity is another weird quantum state of matter showing up at ultracold temperatures, when atoms behave as a crystal and as a superfluid at the same time. Two routes exist. The historical Andreev-Lifshitz-Chester route relies on delocalized vacancies in a crystal lattice which behave quantumly as bosons and switch position by quantum tunnelling like in a Bose-Einstein condensate. The route that actually worked experimen-

tally uses ultracold dipolar gases, where the condensate spontaneously develops a density modulation while keeping a global phase coherence [621].

This state of matter was predicted in 1969 by David J. Thouless (1934-2019, British, 2016 Nobel prize in physics) and, independently, by Alexander Andreev (1939, Russian) and Ilya Mikhailovich Lifshitz (1917-1982, Russian) [622, 623]. A supersolid signature was claimed in 2004 with ^4He at a pressure of about 60 bar and below 170 mK, but the non-classical rotational inertia observed was later traced to a shear stiffening of solid helium and retracted by its own authors in 2012 [624–627]. The first uncontested supersolids came in 2019 with ultracold dipolar quantum gases of erbium and dysprosium in Pisa, Stuttgart and Innsbruck [628, 629]. The related fundamental research is going on in various places in the world like in the USA, Innsbruck (led by Francesca Ferlaino from the University of Innsbruck, IQOQI), Pisa [630], Stuttgart, Warsaw, Geneva, and Paris. It is now possible to create supersolids with ultracold dipolar quantum gases of highly magnetic lanthanide atoms like erbium and dysprosium. The supersolidity effect can be controlled by a magnetic field.

There are no known practical applications of this phenomenon to date although it could lead to new forms of quantum simulation systems like the ones using cold atoms.

3.10.6 Polaritons

Polaritons is a field of quantum physics that is rarely mentioned in the context of quantum technologies. It mostly belongs to fundamental research but could be of interest in various fields such as quantum computing and quantum sensing. Polaritons are quantum quasiparticles in the domain of strong coupling between light and matter. They result from the coupling between photons and an electrical polarization wave.

These waves show up in plasmons (oscillations of free electrons in metals), phonons (oscillations of atoms, especially in crystal structures) and excitons (bound electron-hole pairs generated by photons in semiconductors [631–633]). The polariton naming was created by Joseph John Hopfield (1933, American) in 1958 which at that time concerned polariton excitons. The materials can be atoms gas, massive classical semiconductors, thin films inserted in optical cavities or superconducting Josephson junctions. Hopfield is also known in the field of neural networks in AI with his “Hopfield networks”. He was a laureate of the Nobel prize in physics in 2024.

Excitation photons have a wavelength corresponding to the resonance frequency of the associated medium, often in the visible light or infrared ranges. Polaritons have mixed properties of photons dressed by electronic excitations. They behave like bosons (having an integer spin) that can occupy the same quantum state and operate




Figure 149: Various forms of polaritons. Source: [634], which makes a good inventory of different types of polaritons and their fields of application.

in groups, such as superconducting currents forms with paired electrons named Cooper pairs or Bose-Einstein condensates (BEC).

Depending on the interaction scale, polaritons operate in a semiclassical or quantum regime. In the first case, the electromagnetic field interacts with a macroscopic polarization field. The polariton field then has the properties of a classical field but its elementary quantum is the result of a dipole-photon “wrapping” that can only be described by quantum mechanics. In the second case, the electromagnetic field interacts with a single polarization field quantum that has been isolated in one way or another, such as a superconducting qubit or an exciton in a quantum box. We are then in the quantum regime described by the “Jaynes-Cummings Hamiltonian”. Discrete energy levels and excitation-number sectors exist whatever the coupling strength, so they are not what defines the strong-coupling regime. Strong coupling means that the coherent light-matter coupling rate g exceeds the cavity loss rate κ and the emitter dephasing rate γ , so that the dressed polariton doublets split by $2g\sqrt{n}$ within each excitation-number manifold become spectrally resolvable instead of being washed out by dissipation. Cavity-excited polaritons are generally in the first regime.

In polaritons, semiconductor matter receives photons that excite it. It then emits photons to get out of its excited state, all of this in a very fast iterative cycle, the photons circulating in a closed circuit in the cavity. In practice, electromagnetic and polarization fields co-propagate in the medium in an identical way, notably in polarization and frequency, and with a fixed phase relation (without phase shift or with a 180° phase shift, i.e., π). Polaritons are particularly interesting for generating strong nonlinearities which are searched in photonics [635].

Thanks to the degenerate states in which polaritons can be prepared and to the fact that they interact with each other, polaritons constitute an out-of-equilibrium

quantum fluid called “light quantum fluid”, or “quantum fluids of light” [636], often abusively referred to as “liquid light”.

Polaritons can thus generate surface waves and propagation phenomena typical of quantum fluids such as superfluids.

Polaritons also interact with each other, which is not the case for photons in vacuum [637, 638].

We can experimentally control the spatial distribution of the density, phase and velocity of these fluids of light [639].

There are many variants of polaritons which depend on the nature of the electronic excitation of the matter (Figure 149):

- **Phonon-polaritons** resulting from the coupling between an infrared photon and an optical phonon caused by the mechanical oscillation of two adjacent ions of opposite charge in a crystalline structure. This oscillation produces an oscillating electric dipole moment. This phenomenon was discovered by **Kirill Tolpygo** (1916-1994, Russian) in 1950 and, independently, by **Kun Huang** (1919-2005, Chinese) in 1951. One application of phonon polaritons are thermal emitters and imagers [640].
- **Exciton-polaritons** result from the coupling of a photon with an exciton in a semiconductor cavity. An exciton is a quasiparticle consisting of an electron-hole pair connected by Coulomb forces, generated by excitation photons (Figure 150). The notion of exciton was created by **Yakov Frenkel** (1894-1952, Russian) in 1931. Like all types of polaritons, these have two energy bands: the high and low polariton. It is a general property of the strong coupling regime between electric dipole and electromagnetic field. Here, the level is high when the photon and the semiconductor are excited and in phase, and low when they are in opposite phase. Research is going

on trying to create transistors using polariton-exciton as well as on single quantum control. The “polariton blockade” mechanism allows in principle to manipulate excitonic cavity polaritons at the single quantum scale [641].

- **Surface plasmon polaritons (SPP)** result from coupling surface plasmons and photons. A plasmon is the quantum of a collective electron-density oscillation, whatever the carrier density. A surface plasmon is a coherent electron oscillation occurring at the interface between two different materials, often a metal and a dielectric or between metal and air. A surface plasmon polariton is the hybrid electromagnetic and surface charge-density mode formed when such an oscillation couples to an incident photon, so it is a composite object and not simply a plasmon (Figure 151).

SPPs are used in optical quantum sensors for temperature and for the detection of the concentration of different components by refractivity and then spectroscopy, like in medtechs (detection of various organic molecules and of interactions between proteins), biological analysis (toxins, drugs, additives) or for the detection of gases. The general principle of this instrument is to use a laser diode to illuminate a gold surface at an angle (via a mechanically controllable angle) and to capture the reflected beam with a detector. The gold surface is coated with a specific molecule (“biorecognition element” in the diagram) that tends to associate itself with a molecule that we want to detect (in the liquid phase “flow of analyte”). The molecules detected can be peptides, polypeptides, proteins, enzymes, vitamins, DNA or RNA sequences, or antibodies (in particular for cancers diagnosis). The association modifies the reflectivity of gold and allows the detection of the target molecule.

SPRs (Surface Plasmon Resonance Plasma) can be much more powerful than near-infrared spectroscopy sensors such as those from Scio [642]. They measure the polarized light reflected from a laser diode in terms of intensity, angle, wavelength, phase and polarization. It can be used for various bioanalysis tasks [643].

As in many biological analysis systems (Figure 153), it is possible to create 2D matrices (microarrays) integrating many detection molecules and to detect a lot of components in the sample to be analyzed.

SPRs are commonly marketed by companies such as **Cytiva** (USA), **Carterra** (USA), **Horiba** (Japan), **IBIS Technologies** (Netherlands), **Lifasible** (USA), **Polaritons Technologies** (Switzerland) and **XanTec** (Germany). Horiba’s European research center is located in Palaiseau next to the C2N of the CNRS, Télécom Paris, Thales and the Institut d’Optique. Horiba is specialized in spectrometers and various other optical instruments like for the near-IR photoluminescence characterization of InGaAs/GaAs

quantum dots. They acquired Yvon Jobin, a French optical instruments manufacturer in 1997.




Figure 150: Exciton-polariton. Source: [644].

- **Cavity polaritons** are a variant of the polariton excitons where the photon is trapped in a microcavity, and the exciton is confined in a quantum well. They are made of III-V semiconductors like indium, arsenic and gallium. Photons trapping is often performed using two Bragg mirrors facing each other to create an optical cavity using layers of dielectrics to reflect light very efficiently over a stopband centered on the cavity wavelength. These mirrors are fabricated from molecular beam epitaxy allowing coherent crystal growth on a gallium arsenide (GaAs) crystal substrate. The result is monocrystalline and can contain more than a hundred layers of different alloys, with thicknesses ranging from 5 nm to 50 nm, controlled to the nearest atomic monolayer [646]. These microcavity polaritons were discovered in 1992 by Claude Weisbuch (France) [647].
- **Intersubband-polaritons** result from the coupling of an infrared or terahertz photon with an intersubband excitation. They can be used to create infrared detectors.
- And then **Bragg-polaritons** (Braggoritons), **plexcitons** (plasmons + excitons), **magnon polaritons** (magnon, spin waves in ferromagnetic materials + photons) and **similaritons** (amplified photons in optical fibers).

In short, all these “*-ons” are the result of the interaction between photons and different forms of matter, noticeably electrons. What does this have to do with quantum computing? Polaritons are used in various optical devices related to photon qubits, including photon transport and single photon detectors. They could eventually allow the creation of photon qubits that can interact with each other. This is what emerged from an




Figure 151: Surface-plasmon polariton phenomenon. Source: [Wikipedia](#).




Figure 152: Visualizing a skyrmion. Source: [\[645\]](#).




Figure 153: Surface plasmon resonance plasma. Source: [Surface Plasmon Resonance \(SPR\)](#) by Lifeasible.

MIT and Harvard publication by Vladan Vuletić and Mikhail Lukin in 2018 which observed a three-photon bound-state component in light propagating through an ultracold atomic gas under Rydberg blockade, the photons being dressed into interacting “Rydberg polaritons” [\[648\]](#). This established strong photon-photon interaction in one specific Rydberg EIT medium. It did not demonstrate a scalable photonic gate and no gate fidelity was reported. Another research project in Singapore uses polariton excitons to create photon qubits with the particularity of being able to operate at room temperature, using single-qubit gates and $\sqrt{SWAP}$ two-qubits gates [\[649\]](#).

Microcavities polaritons can be used to create quantum simulators [\[650, 651\]](#). They are implanted in III-V semiconductor structures as 2D arrays. One field of application is the simulation of gravitational structures such as a Hawking radiation on the horizon of a black hole. And why not, to simulate the operation of a dilution refrigerator associating helium 3 and 4 at very low temperature.

Polaritons are also the field of topological behaviors of matter and are perhaps an alternative way to the Majorana fermions to create error corrected qubits. These

are longer term pathways than the qubit technologies studied in this book, but worthy of interest.

Other applications, already mentioned, target the very diverse field of quantum sensing, including optomechanical systems [\[652\]](#).

In France, polaritons are the specialty of Cristiano Ciuti (UPC MPQ), **Elisabeth Giacobino** (CNRS LKB) [\[653\]](#), **Jacqueline Bloch** (CNRS C2N), **Alberto Bramati** (ENS LKB), **Alberto Amo** (PhLAM CNRS Lille), **Le Si Dang** and **Maxime Richard** (CNRS Majulab Singapore).

3.10.7 Magnons

Quantum matter also includes magnons, a category of quasiparticles that take the form of quantized spin waves in magnetic materials, usually crystalline lattices. Magnons were conceptualized by **Felix Bloch** in 1930 and experimentally detected in 1957 by **Bertram Brockhouse** (1918-2003, Canadian), who was awarded the Nobel prize in physics in 1994 for neutron spectroscopy. These objects which behave as bosons could be used in quantum information systems.

Current physics experiments are done at the control low-level like with controlling these magnons with mi-

crowaves [654] or measured with superconducting qubits [655, 656]. Magnons can also be used at low temperature to create some topological materials [657] and even for some species of SiC-based spin qubit control [658], and hybrid quantum sensors [659, 660].

3.10.8 Skyrmions

Order is not restricted to the periodic atomic array of a crystal and can also be associated with magnetic order in a solid where spins align parallel to each other in ferromagnets and antiparallel in antiferromagnets. Skyrmions are more complex magnetic nanostructures that form mesoscopic magnetic vortex with particle-like properties [661].

How to distinguish magnons and skyrmions which are both magnetic quasiparticles? Magnons are quantized dynamic magnetic excitations that travel through magnetic materials while skyrmions are static. Magnon and skyrmion are two phenomena which can be combined [662]. The skyrmion naming comes from **Tony Hilton Royle Skyrme** (1922-1987) who formulated in 1961 a nonlinear field theory of massless pions in which particles can be represented by 3D topological solitons, *aka* hopfions [663, 664], which can lead to the creation of time crystals in nematic liquid crystals [665].

Skyrmions existence in magnetic materials was predicted in 1989 by Bogdanov et al. [666]. In 2009, **Sebastian Mühlbauer** and colleagues discovered a skyrmion lattice in MnSi crystals with neutron scattering at the FRM II reactor in Munich [667].

Skyrmions have a topological counterpart, anti-skyrmions [668], with an opposite topological charge and a distinct internal structure. Their nonzero topological charge gives them metastability, since no continuous deformation can unwind them into the uniform state in the continuum limit, which translates into a real energy barrier. The protection is however relative and not absolute. Real textures are nucleated and annihilated at sample boundaries and at lattice defects, by magnetic fields, by spin-polarized currents and by thermal activation, and the discreteness of the lattice itself allows the topological charge to change. There are also other textures such as hedgehogs, instantons and merons, which are distinct objects belonging to different dimensionalities and order parameters rather than interchangeable names. A bimeron is a bound pair of merons and is the in-plane magnetized counterpart of a skyrmion [669]. In MnSi, the skyrmion lattice only exists in a small pocket of the phase diagram just below the 29.5K magnetic ordering temperature.

Researchers in Japan and South Korea implemented real-space imaging of a two-dimensional hexagonally arranged skyrmion lattices spaced by 90 nm in a thin film of $\text{Fe}_{0.5}\text{Co}_{0.5}\text{Si}$ and exposed to a magnetic field of 50–70mT, using Lorentz transmission electron microscopy [645].

This helicoidal structure can also be 3D and create superposition of various magnetic skyrmion states (Figure 152). This could lead to the creation of new ultra-high-density memories [670] particularly with the room-temperature Néel skyrmions that can be made with thin-film systems [671, 672] and with cold atoms [673], to in-memory processing architectures [674, 675], neuromorphic processors [676], to QRNGs [677], in low-power spintronic applications [678, 679] and for the creation of isolators to be used in superconducting read-out chains [680]. Skyrmions could also a new breed of qubits with magnetic nano disks bounded by electrical contacts, where static electric and magnetic fields control the skyrmions quantized energy levels corresponding to their helicity. A way to entangle them must then be found [681, 682]! Skyrmions can also be controlled by mechanical systems [683].

3.10.9 Topological matter

The very concept of topological quantum states leading to topological matter was discovered with a specific insulating phenomenon that can be explained by the **quantum Hall effect**, with electrons moving through a strong magnetic field and accumulating in some parts of the material depending on its shape. This electron conductivity is quantized, as discovered in 1980 by **Klaus von Klitzing** (Germany) who was awarded the Nobel prize in physics in 1985. This “integer” quantum Hall effect was later completed by the discovery of the fractional quantum Hall effect by Tsui et al. in 1982 in two-dimensional electron systems in semiconductor devices, followed by the theoretical discovery of the entangled gapped quantum spin-liquid state of integer-spin “quantum spin chains” by **Frederick Duncan Michael Haldane** in 1983, who was awarded the Nobel prize in physics in 2016 along with **David J. Thouless** and **J. Michael Kosterlitz** [684, 685].

In 2005, **Eugene Mele** and **Charles Kane** predicted that topological insulation could happen in graphene sheet submitted to strong spin-orbit coupling creating the quantum Hall effect without any applied magnetic field [686]. This phenomenon is named the “quantum spin Hall effect” and relates to the Kane-Mele invariant [687, 688]. It was demonstrated to occur in wafers of mercury telluride, predicted in 2006 by **Shou-Cheng Zhang** et al. from Stanford University and measured in 2007 by Markus König in **Laurens Molenkamp**’s group at Würzburg [689]. The following year, the first 3D topological insulator, bismuth antimonide, was identified by **Zahid Hasan**’s group from Princeton [690].

Since then, high-throughput screenings have shown that topological materials are anything but rare. A 2019 catalogue covering 26,938 stoichiometric materials from the Inorganic Crystal Structure Database identified 3,307 topological insulators and 4,078 topological semimetals, meaning that more than 27% of all known materials host some topological band structure, a figure later con-

firmed on the full set of nonmagnetic stoichiometric materials [691–694] (Figure 154). A French American research team devised in 2020 a machine learning model to detect such topological insulators out of an initial database of 4,009 candidates [695]. Again, spintronics are a potential use case of topological insulators to create power-saving electronics where the on/off of a bit would be an electron spin instead of the on/off path of an electron stream.

In topology, the invariant characterizing a phase depends on its symmetry class and on its dimensionality, and there is no single universal one. A 2D system with broken time reversal symmetry, as in the integer quantum Hall effect, is labeled by an integer **Chern number**. A 1D chiral system such as the Su-Schrieffer-Heeger chain is labeled by a winding number. A time reversal invariant topological insulator is labeled by a $\mathbb{Z}_2$ index, and crystalline topological phases require yet other invariants. These are not interchangeable. What they share is that the invariant can only change when the bulk energy gap closes, which is precisely what happens at a topological quantum phase transition. These are other various physics concepts to consider, way beyond what I can do at this point in my quantum journey [696–698].

It is interesting to note that some materials can showcase 3D topological behavior at ambient temperature, like bismuth-selenide (Bi_2Se_3). It is a semiconductor and a thermoelectric material that has a topological insulator ground-state. Bi_2Se_3 nanostructures have been investigated for radiosensitization in targeted cancer therapy and for X-ray and mammography detection [699]. These proposals rest on the high atomic number of bismuth and on the semiconducting and thermoelectric properties of the compound. Whether the topological surface states are functionally required has not been demonstrated. You can also potentially build **magnetic monopoles** quasiparticles, breaking the convention that magnetism always shows up with dipoles [700].

You are certainly willing to “visualize” the different types of topological materials identified. I found this nice and highly detailed table showing their great diversity in a review paper, in Figure 155.

Topological matter can have several applications related to light-matter interactions in the Terahertz regime. It can help create waveguides, optical isolators and diodes that are more resistant to their environment perturbations in the recent field of **topological photonics** which is related to polaritons [702–704].

We even have **topological lasers**, which can for example consolidate multiple sources in a coherent way, leading to even more powerful lasers, using a topological insulator vertical-cavity surface-emitting array (VCSEL) [705–707].

Then of course, one key application of topological matter is topological qubits, often associated with Majorana fermions sought after by Microsoft. But topological

qubits are way more diverse with many competing definitions and architectures. For example, you also can count with Fibonacci anyons [708].

3.10.10 Time crystals

Time crystals are beasts we hear a lot about since mid-2021, when Google announced it had created such artefact in its Sycamore processor [709, 710]. It shed some light on this weird phenomenon that was devised in a 2012 paper by Frank Wilczek from the MIT (and 2004 Nobel prize in physics awardee) and by another paper by him and Alfred Shapere from the University of Kentucky [711, 712].

This thing is somewhat linked to the history of the search for a perpetuum movement, an isolated object supposed to keep in motion indefinitely. It was dismissed by the French Academy of Science in 1775 due to the limits of friction and, later, to the second law of thermodynamics [713, 714].

In classical crystals, atoms are periodically arranged in space structured according to one of the 230 structures already described. A discrete time crystal is the analogue along the time direction, and not something arranged in space and time at once. It is a periodically driven many-body system whose observables respond at a subharmonic of the drive, most often at half its frequency, robustly against changes of the drive parameters, which spontaneously breaks the discrete time translation symmetry of that drive [715, 716]. No spatial crystalline order is required for this. The scientific description of the phenomenon is the less explicit “spontaneous time symmetry breaking”. There’s even an acronym for this, TTSB which means time translation symmetry breaking. There are also continuous time crystals that were observed first in 2022 in Germany [717–719].




Figure 156: Time crystal oscillations over time. Source: [720].

A time crystal is a nonequilibrium dynamical phase of matter and not a perpetual motion machine. Its subharmonic oscillation is not motion without energy. The drive keeps supplying energy, driven and open realizations do dissipate, and many of them need that exchange




Figure 154: A classification of topological matter. Source: [Research Lines - Theory of Topological Matter](#) by Adolfo Grushin, CNRS.

Phases	System	Sym.	Edge/Surface	Band	FS	DOS
Integer Quantum Hall Insulator (1980)	Heterostructure with 2DEG and Disorder		 2D bulk State			
Fractional Quantum Hall Insulator (1982)	Clean Heterostructure With 2DEG		 2D bulk State			
Quantum Spin Hall Insulator (2005)	Clean Heterostructure with band inversion		 2D bulk State			
2D Dirac Semimetal (2005)	2D Materials (Graphene) With Intrinsic SOC					
3D Topological Insulator (2007)	3D Insulator Materials					
Topological Magnetic Insulator (2010)	Topological Insulators with magnetic doping					
Topological Kondo Insulator (2010)	Strong Correlated System with Heavy Fermions					
Topological Superconducting Insulator (2010)	Topological Insulators with Superconducting doping					
Topological Crystalline Insulator (2010)	Bulk Insulators					
Dirac Semimetals (2012)	Bulk Semimetals					

Phases	System	Sym.	Edge/Surface	Band	FS	DOS
Dirac Nodal Line Semimetals (2015)	Bulk Semimetals					
Weyl Nodal Line Semimetals (2016)	Bulk Semimetals					
Type I Weyl Semimetals (2015)	Bulk Semimetals					
Type II Weyl Semimetals (2015)	Bulk Semimetals					

Legends:

Symmetries:

Time-Reversal Inversion
Time-Reversal Breaking Time-reversal + Inversion

Temperature: Low Temperature Room Temperature

Field: B-field Spin-orbit Coupling

Rotational **Mirror**
Glide

DCS

Figure 1: Tabular illustration of different 2D and 3D quantum/topological families along with the corresponding material systems, symmetries that help to create the phase, topological edge or surface states in real and momentum space, band structures, Fermi surfaces (FS), and density of states (DOS). All topological phases are also categorized according to topological insulators (black text) and semimetals (green text). In Legends, TRS, IS, I+TRS, and DCS represents time-reversal, inversion, inversion + time-reversal, and discrete crystal symmetry, respectively.

Figure 155: A table with a classification of various topological materials in 2D and 3D and indicating time reversal and operating temperature. Source: [701].

with a bath or an active stabilization mechanism, such as many-body localization or prethermalization, to keep the system from heating up toward infinite temperature. What is robust is the periodicity of the response, not the absence of energy flow. Some real time crystals were first observed in lab experiments, starting in 2017 with a periodically driven linear chain of ten trapped ytterbium ions (which by the way, shows some signal damping, in Figure 156) [720] and, the same week, with an ensemble of about one million NV centers in diamond at room temperature at Harvard [721]. It can also happen with some continuous change of spin for some particles, when

the change periods is up to 100 times longer than the system drive period. It was tested in 2021 by a QuTech team in The Netherlands using controllable ¹³C nuclear spins in diamond structures [722]. Things get complicated when you learn that time crystals have also been experimented with superconducting qubits like with the Google 2021 experiments and other subsequent ones with a continuous line of 57 qubits in a 65 qubits IBM QPU [723]. How could a series of connected superconducting qubits become a “crystal” per se?

The word crystal here refers to the spontaneous breaking of a time translation symmetry in the measured observables and not to the material construction of the device. A chain of coupled superconducting qubits is a many-body spin system, and whether it is built out of aluminum, aluminum oxide, niobium and titanium is irrelevant to the temporal order being observed, in the same way that a spatial crystal needs no clock.

So why all this fuss around time crystals and how could they become useful? Some think they may be useful to create some form of quantum memory. Others mention ways to improve NMR spectroscopy, AMO-based quantum simulation (AMO standing for Atomic, Molecular, and Optical physics), to stabilize Schrodinger-cat states against local perturbations, enhance metrological bandwidth while maintaining sensitivity like with the measurement of magnetic fields, as a frequency standard or for beyond-SQL (standard quantum limit) quantum sensing [147].

3.10.11 Quantum batteries

The idea of quantum batteries to improve energy storage and disposal was created in 2012-2013 by Robert Alicki (University of Gdansk) and Mark Fannes (Universiteit Leuven). This fundamental research dealt with finding how much work could be stored and extracted from quantum systems thanks to the use of quantum entanglement, *aka* quantum batteries [724].

Quantum batteries could store energy in high energy states of quantum objects and release it efficiently. The advantage discussed in the literature is a *superextensive charging power*, meaning a total charging power growing faster than linearly with the number of quantum cells or spins, and it only holds under explicitly matched control resources. The stored energy itself stays extensive and scales with the number of cells just like in a classical battery, and no generic gain in energy density or in retention time follows from entanglement [725–727]. Quantum batteries belong to the emerging field of quantum energy science [728].

Theoretical work. In the cavity Dicke model studied by [729], a quantum battery made of N two level quantum systems coupled to a shared cavity mode has an effective collective coupling scaling as $\sqrt{N}$, so the charging time scales as $1/\sqrt{N}$ instead of staying independent of N as it would for N cells charged in parallel. This is a model-specific collective-coupling result under its own assumptions and not a general law of quantum batteries. Any fair comparison has to fix the total driving strength, the available energy and the coupling resources across the collective and the parallel protocols, otherwise the $\sqrt{N}$ merely reflects a stronger field. This superabsorption property was demonstrated with cold-atoms [730] and organic-molecules in cavity [731]. A so-called Dicke model was developed to explain superradiance, where atoms emit light collectively and more intensely than individually, in the quantum battery discharge. N two-

level systems like atoms or qubits are placed in a cavity and interact coherently with a shared photon field. The collective interaction allows for superabsorption and ultrafast charging. However, it seems that this scaling comes from the stronger cavity electric or magnetic field and not to some (scalable) quantum effect [197, 732].

The work on quantum batteries is still very theoretical and quite far from practical batteries [733, 734]. The main benefit of these quantum batteries seems to be fast charging, with the caveat of fast discharging, which is quite inconvenient [735]. There are no available data on potential energy density in a real documented manner with full-stack product packaging, one of the main showstoppers for various use cases like for long distance electric vehicles or aerial vehicles. So, you're far from buying your next Tesla equipped with a 1,000-mile range quantum battery [736], particularly given most quantum batteries experiments run in ultra-cold environments to avoid quantum decoherence [737]. So, what do we have in-store here? Mainly scientific work with very low TRLs [738, 739].

First prototypes. On the practical side, quantum battery prototypes use organic dye molecules, photonic cavities, NV centers in diamond, photons [740], and quantum spin chains [741, 742]. The reported figures deserve care, since absorbed energy, internal energy, extractable work (the ergotropy), charging power, retention time before self-discharge and full-system energy input are six different quantities and most papers report only the first. Most quantum battery experiments deposit nanojoule to microjoule scale excitation energies that survive nanoseconds to microseconds, with no measured extractable work, no discharge protocol and no efficiency figure covering the laser and the cryogenics. They currently need room-sized laboratory setups for storing a couple microjoules. Thus, scaling up remains a challenge. This is a perfect example of the gap that exists between fundamental research and engineering.




Figure 157: A photonic quantum battery storing the energy of one photon. This is the early stage before scaling! Source: [740].

Scientists from Australia and Italy published work in 2022 on an **organic battery** with fast charging using superextensive scaling of absorption [731]. It was based on a thin active layer of a low-mass molecular semiconductor named LFO (Lumogen F-orange) that

was dispersed into a polymer matrix that is sandwiched between two dielectrics made of 8 and 10 pairs of Bragg mirrors, creating a microcavity. The battery cell was then controlled by a laser at about 500 nm, in the blue-green range, a noncollinear optical parametric amplifier, beam splitters and delay lines and a detector. In a word, we could say it is a “light” battery, absorbing energy as light, and rendering it as light, in a different wavelength. Like in many other papers of this kind, it is quite difficult to infer the practicality of these quantum batteries (Figure 158).

Other research work:

- A team from Italy and Korea proposed in 2022 to use micromasers to store energy with microwave photons [744].
- Other scientists from Australia are looking at ways to store energy in light-induced spin state trapping in spin crossover materials [745, 746].
- Researchers from various laboratories in China are also working on quantum batteries fundamental research [747, 748], including batteries using atom nuclei [749].
- A Korean American German Singaporean team described in 2022 how quantum batteries are isolated quantum systems undergoing unitary charging protocols (unitary in the mathematical sense). The described system was based on an Otto engine which can serve as an engine and as a refrigerator [750].

Room-temperature prototype. In 2025, based on James Quach’s work in 2022 [731], a room-temperature quantum battery prototype was presented. It used organic microcavities and polariton states with energy storage lasting microseconds, the next goal being to reach seconds. It used layered materials and reflective cavities to trap laser light and maintain coherence without extreme cooling [751]. A full-stack energy analysis is not provided in the study. Ideally, it should account for the poor laser energy efficiency that is usually between 15% and 30% for the diode lasers used in such settings (Figure 160).

Overselling. If researchers are not overselling it, the news media are doing it, touting “*batteries with one million miles autonomy*” [752]. It came from another paper, authored by Canadian scientists and an engineer from Tesla which proposed in 2019 an improved Li-Ion battery that could last 1.5 million miles over its lifespan but, of course, not with a single recharge [753]. And it was not even a quantum battery. There is still some practical interest to develop quantum batteries at the quantum level. It can help manage energy transfers

and storage with qubit drive like with superconducting qubits [754, 755].

Quantum inspired batteries? In some other work from US and Japanese researchers, we are closer to classical battery designs. It is about using lithium-doped samarium nickelate, a quantum crystalline material with strongly correlated electron systems [743] (Figure 159). Lithium ions are usually the main compound of batteries electrolytes. The quantum crystal structure improves the conduction of these ions that could also be sodium ions. It could enable better electrolytes but another effect of the structure where additional electron modifies the material conductivity could be used in neuromorphic synapses for storing neural networks connections weights.

Other research work deal with microscopic batteries which don’t seem to be useful for energy storage [756, 757]. They can help better understand the thermodynamics of qubits manipulation and provide innovative insights on how to fight decoherence and noise [216].

Higher TRLs can be found with rather classical batteries that would use topological semi-metallic porous carbon materials as potential more efficient anodes for Li-Ion, sodium-ion and potassium-ion batteries. Other topological materials could be useful for supercapacitors and to create more efficient catalysts for water electrolysis, with the production of hydrogen in sight coming from renewable originated electricity [758]. Coherence-based quantum capacitors are a distinct and much more speculative proposal [759].

PioniQ Technologies

France, 2024



PioniQ Technologies spun out from ESPCI (École Supérieure de Physique et de Chimie Industrielles de la Ville de Paris) which is developing energy storage solutions based on a ceramic perovskite-derived material dubbed “h-Bar”. As of 2026, the company describes its technology as a new class of Li-free, Co-free, Ni-free solid-state energy storage devices “*powered by quantum matter phase transitions*”. This is a company-reported claim. No capacity, power density, cycle life, round-trip efficiency, operating temperature range, TRL or independent validation has been published to date, and the mechanism has not been described in a peer-reviewed paper. Clément Barraud is their CEO and Quantonation is a seed investor in the company.

QpiVolta (2024, India) is developing a solid-state battery which the company describes as relying on a quantum matter phenomenon. As with PioniQ, this is a company-reported claim as of 2026, with no published capacity, power, cycle life or efficiency figures and no independent validation.




Figure 158: The first quantum battery concept experimented by James Quach in Australia in 2022. Source: [731].




Figure 159: Lithium-doped samarium nickelate quantum battery. Source: [743].




Figure 160: The first room-temperature prototype quantum battery was created in 2025, with energy storage lasting microseconds. It is using a multi-layered microcavity design using various materials like an absorber molecule, copper phthalocyanine (CuPc), as well as silver and fullerene (C₆₀). Source: [751].



3.10.12 Bibliography and notes

These are the bibliographical references and notes for the **Quantum matter** subsection.

- [147] [Colloquium: Quantum and Classical Discrete Time Crystals](#), by Michael P. Zaletel, Mikhail Lukin, Christopher Monroe, Chetan Nayak, Frank Wilczek, and Norman Y. Yao, arXiv, May 2023 (29 pages).
- [197] [Reply to the comment on “High-Power Collective Charging of a Solid-State Quantum Battery” by Haowei Xu and Ju Li](#), by D. Ferraro, M. Campisi, G. M. Andolina, V. Pellegrini, and M. Polini, arXiv, November 2024 (2 pages).
- [216] [Experimental analysis of energy transfers between a quantum emitter and light fields](#), by Ilse Maillette de Buy Wenniger, Alexia Auffèves, Pascale Senellart, et al., arXiv, June 2023 (18 pages).
- [509] [Exploring Quantum Materials & Applications: A Review](#), by Rajat Kumar Goyal, arXiv, April 2024 (60 pages).
- [510] [CarboQuant](#).
- [511] [Basic Research Needs for Quantum Materials](#), by Department of Energy, DoE, 2016 (4 pages).
- [512] [Quantum Materials for Energy Relevant Technology](#), by the DoE Office of Science, 2016 (170 pages).
- [513] [Quantum materials: Where many paths meet](#), by Philip Ball, MRS Bulletin, October 2017 (8 pages).
- [514] [Magic-Angle Multilayer Graphene: A Robust Family of Moiré Superconductors](#), by Jeong Min Park, Pablo Jarillo-Herrero, et al., arXiv, December 2021 (15 pages).
- [515] [Magic-angle twisted symmetric trilayer graphene as a topological heavy-fermion problem](#), by Jiabin Yu, PRB, July 2023 (65 pages) , (preprint on [osti.gov](#)).
- [516] [Evidence for Flat Band Dirac Superconductor Originating from Quantum Geometry](#), by Haidong Tian, Marc W. Bockrath, et al., arXiv, December 2021 (15 pages).
- [517] [The 2021 Quantum Materials Roadmap](#), by Feliciano Giustino, Stephan Roche, et al., arXiv, February 2021 (93 pages).
- [518] [Introduction to Quantum Materials](#), by Leon Balents, 2018 (51 slides).
- [519] [Quantum Matter Overview](#), by Melanie Swan et al., April 2022 (23 pages).

- [520] [On demand defining high-quality, blue-light-active ZnSe colloidal quantum wires](#), by Yi Li et al., National Review Science, April 2022.
- [521] [Relativistic Phases in Condensed Matter](#), by Thibaud Louvet, 2018 (165 pages).
- [522] [After 85-year search, massless particle with promise for next-generation electronics discovered](#), by Morgan Kelly, 2015.
- [523] [Topological Semimetals](#), by Andreas P. Schnyder, 2020 (32 pages).
- [524] [Imaging Orbital Vortex Lines in Three-Dimensional Momentum Space](#), by T. Figgemeier et al., Physical Review X, February 2025 (10 pages).
- [525] [Realization of a Laughlin quasiparticle interferometer: Observation of fractional statistics](#), by F. E. Camino, Wei Zhou, and V. J. Goldman, arXiv, September 2005 (25 pages).
- [526] [On the theory of identical particles](#), by J. M. Leinaas and J. Myrheim, Il Nuovo Cimento B, January 1977 (23 pages) .
- [527] [Fractional statistics in anyon collisions](#), by H. Bartolomei, G. Feve, et al., Science, April 2020 (9 pages).
- [528] [Direct observation of anyonic braiding statistics](#), by J. Nakamura, S. Liang, G. C. Gardner, and M. J. Manfra, Nature Physics, September 2020 (20 pages).
- [529] [A Field Guide to Spin Liquids](#), by Johannes Knolle and Roderich Moessner, arXiv, April 2018 (17 pages).
- [530] [Scale-invariant magnetic anisotropy in RuCl₃ at high magnetic fields](#), by K. A. Modic, Arkady Shekhter, et al., arXiv, May 2020 (32 pages).
- [531] [Resonating valence bonds: a new kind of insulator?](#), by P. W. Anderson, Materials Research Bulletin, 1973 .
- [532] [Quantum Spin Liquids](#), by C. Broholm, R. J. Cava, S. A. Kivelson, D. G. Nocera, M. R. Norman, and T. Senthil, arXiv, May 2019 (21 pages).
- [533] [Quantum Glasses – a review](#), by L. F. Cugliandolo and M. Mueller, arXiv, September 2022 (24 pages).
- [534] [Topological Insulators](#), by M. Z. Hasan and C. L. Kane, arXiv, November 2010 (23 pages).
- [535] [Topological Insulators](#), by Yoichi Ando, arXiv, July 2023 (10 pages).
- [536] [A Beginner's Guide to Topological Materials](#), by Charles Q. Choi, IEEE Spectrum, June 2021.
- [537] [Imaging the Wigner Crystal of Electrons in One Dimension](#), by Ilanit Shapir, Shahal Ilani, et al., arXiv, March 2018 (38 pages).
- [538] [Observation of Wigner crystal of electrons in a monolayer semiconductor](#), by T. Smoleński, A. Imamoglu, et al., arXiv, October 2020 (7 pages).
- [539] [Nanomatériaux et nanotechnologies : quel nanomonde pour le futur?](#), by Pierre Rabut, 2018 (23 pages).
- [540] [Advances in Quantum-Dot-Based Displays](#), by Yu-Ming Huang, 2020 (29 pages).
- [541] [Quantum dots and their potential impact on lighting and display applications](#), by Paul W. Brazis, 2019 (18 pages).
- [542] [Crystal Systems and Space Groups](#), by Paul D. Boyle et al., 2014 (44 slides).
- [543] [Chirality and the next revolution in quantum devices](#), by César Tomé López, May 2022.
- [544] [Topology and Chirality](#), by Claudia Felser and Johannes Gooth, arXiv, May 2022 (27 pages).
- [545] [Magnetic Group Table, Part 2 Tables of Magnetic Groups](#), by Daniel B. Livin, 2014.
- [546] [Exhaustive constructions of effective models in 1651 magnetic space groups](#), by Feng Tang and Xiangang Wan, arXiv, May 2021 (25 pages).
- [547] [Structure and Topology of Band Structures in the 1651 Magnetic Space Groups](#), by Haruki Watanabe, Hoi Chun Po, and Ashvin Vishwanath, arXiv, August 2018 (43 pages).
- [548] [Experimental demonstration of input-output indefiniteness in a single quantum device](#), by Yu Guo, Guang-Can Guo, Giulio Chiribella, et al., arXiv, April 2024 (18 pages).
- [549] [Experimental superposition of time directions](#), by Teodor Strömberg, Philip Walther, et al., arXiv, May 2024 (15 pages).
- [550] [Superconductivity and Electronic Structure](#), by Alexander Kordyuk, 2018 (145 slides).
- [551] [Theory of Superconductivity](#), by J. Bardeen, L. N. Cooper, and J. R. Schrieffer, Physical Review, December 1957 (30 pages).
- [552] [50 Years of BCS Theory “A Family Tree” Ancestors BCS Descendants](#), by Douglas James Scalapino, John Rowell, and Gordon Baym, 2007 (52 slides).
- [553] [The rise of superconductors](#), by P.J. Ford et al., 2005 (224 pages).
- [554] [Microscopic Derivation of the Ginzburg-Landau Equations in the Theory of Superconductivity](#), by L. P. Gorkov, Soviet Physics JETP, 1959 (5 pages).
- [555] [Ein neuer Effekt bei Eintritt der Supraleitfähigkeit](#), by W. Meissner and R. Ochsenfeld, Naturwissenschaften, November 1933 (2 pages) .
- [556] [Theory of Superconductivity](#), by Carsten Timm, February 2022 (150 pages).
- [557] [Type-III Superconductivity](#), by M. Cristina Diamantini et al., March 2023 (7 pages).
- [558] [Topological gauge theory of vortices in type-III superconductors](#), by M. C. Diamantini et al., Physical Review B, September 2024  (preprint on [arXiv](#)).
- [559] [High-Tc superconductivity above 130 K in cubic MH4 compounds at ambient pressure](#), by Xinxin Li, Miguel A. L. Marques, et al., arXiv, November 2025 (9 pages).
- [560] [Exploration of new superconductors and functional materials and fabrication of superconducting tapes and wires of iron pnictides](#), by Hideo Hosono, Satoru Fujitsu, et al., arXiv, May 2015 (141 pages).
- [561] [3DSC - A New Dataset of Superconductors Including Crystal Structures](#), by Timo Sommer, Roland Willa, Jörg Schmalian, and Pascal Friederich, arXiv, December 2022 (15 pages).
- [562] [Helium Reclaiming Magnetic Resonance Imagers](#), by Dan Hazen, 2015 (5 pages).
- [563] [High power density superconducting rotating machines—development status and technology roadmap](#), by Kiruba S Haran, Ernst Wolfgang Stautner, et al., Superconductor Science and Technology, November 2017.
- [564] [Design of MW-Class Ship Propulsion Motors for US Navy by AMSC](#), by Swarn S. Kalsi, 2019 (50 slides).
- [565] [Prospects for innovative power grid technologies, Final report](#), by Compass Lexecon, June 2024 (136 slides).
- [566] [Long Island HTS Power Cable](#), by Department of Energy, Department of Energy, 2008 (2 pages).
- [567] [The Munich SuperLink project](#), by Werner Prusseit et al., 2021 (6 pages).
- [568] [This is what high-voltage superconductors can do](#), by SWM, November 2023.
- [569] [Josephson Junctions in the Age of Quantum Discovery](#), by Hyunseong Kim, Long B. Nguyen, et al., arXiv, May 2025 (39 pages).
- [570] [Multi-purpose quantum laboratories from superconducting circuits](#), by Arpit Arora, Emily M. Been, William Munizzi, Joel Wang, Taylor L. Patti, Aaron Chou, and Prineha Narang, arXiv, August 2026 (19 pages).
- [571] [Superconducting Qubit Based on Twisted Cuprate Van der Waals Heterostructures](#), by Valentina Brosco et al., Physical Review Letters, January 2024  (preprint on [arXiv](#)).
- [572] [Dually tunable YBCO coplanar waveguide resonators based on helium-ion-generated Josephson inductances](#), by Kenny

- Fohmann, Daniel Bothner, et al., arXiv, May 2026 (28 pages).
- [573] [Rapid single flux quantum T-flip flop operating up to 770 GHz](#), by Wei Chen, Alexander V. Rylyakov, Vijay Patel, James E. Lukens, and Konstantin K. Likharev, IEEE Transactions on Applied Superconductivity, June 1999 (4 pages) .
- [574] [Superconductor ICs: the 100-GHz second generation](#), by Darren Brock et al., 2000 (7 pages).
- [575] [A Power Efficient Artificial Neuron Using Superconducting Nanowires](#), by Emily Toomey, arXiv, 2019 (17 pages).
- [576] [Possible High Tc Superconductivity in the Ba-La-Cu-O System](#), by J. G. Bednorz and K. A. Muller, Zeitschrift fur Physik B, June 1986 (3 pages) .
- [577] [Superconductivity at 93 K in a New Mixed-Phase Y-Ba-Cu-O Compound System at Ambient Pressure](#), by M. K. Wu, C. W. Chu, et al., Physical Review Letters, March 1987 (4 pages).
- [578] [Superconductivity at 39 K in magnesium diboride](#), by Jun Nagamatsu, Jun Akimitsu, et al., Nature, March 2001  (preprint on [ResearchGate](#)).
- [579] [Iron-Based Layered Superconductor La\[O_{1-x}F_x\]FeAs \(x = 0.05-0.12\) with T_c = 26 K](#), by Yoichi Kamihara, Takumi Watanabe, Masahiro Hirano, and Hideo Hosono, Journal of the American Chemical Society, February 2008 (2 pages) .
- [580] [Iron-based superconductors: Teenage, complex, challenging](#), by Qimiao Si and Nigel E. Hussey, Physics Today, May 2023  (preprint on [osti.gov](#)).
- [581] [Superconductivity and its applications to society](#), by Amalia Ballarino, INFIERI Summer School 2025, September 2025 (57 slides).
- [582] [Superconductivity at 250 K in lanthanum hydride under high pressures](#), by A. P. Drozdov, M. I. Erements, et al., Nature, May 2019 (24 pages).
- [583] [Quantum Crystal Structure in the 250 K Superconducting Lanthanum Hydride](#), by Ion Errea, Francesco Mauri, José A. Flores-Livas, et al., arXiv, July 2019 (20 pages).
- [584] [Evidence for Superconductivity above 260 K in Lanthanum Superhydride at Megabar Pressures](#), by Maddury Somayazulu, Russell J. Hemley, et al., Physical Review Letters, January 2019 (6 pages)  (preprint on [arXiv](#)).
- [585] [Anomalous High-Temperature Superconductivity in YH₆](#), by Ivan A. Troyan, arXiv, 2019 (36 pages).
- [586] [Synchrotron infrared spectroscopic evidence of the probable transition to metal hydrogen](#), by Paul Loubeyre, Florent Occelli, and Paul Dumas, Nature, January 2020 (5 pages) .
- [587] [Here comes metallic hydrogen - at last!](#), by Jean-Baptiste Veyrieras, May 2020.
- [588] [Room-temperature superconductivity in a carbonaceous sulfur hydride](#), by Elliot Snider, Nature, October 2020 (14 pages).
- [589] [High-Temperature Superconducting Phase in Clathrate Calcium Hydride CaH₆ up to 215 K at a Pressure of 172 GPa](#), by Liang Ma, Yanming Ma, et al., Physical Review Letters, April 2022.
- [590] [Single-photon detection using high-temperature superconductors](#), by I. Charaev, K. K. Berggren, et al., arXiv, August 2022 (8 pages).
- [591] [Novel approach to Room Temperature Superconductivity problem](#), by Ivan Timokhin et al., arXiv, 2020 (4 pages).
- [592] [Scientist claimed to obtain superconductivity at ambient temperature](#), by Amit Malewar, August 2018.
- [593] [A Superconductor Scandal? Scientists Question a Nobel Prize-Worthy Claim](#), by Shannon Hall, Scientific American, August 2018 .
- [594] [‘A very disturbing picture’: another retraction imminent for controversial physicist](#), by Dan Garisto, Nature, July 2023 .
- [595] [Exclusive: official investigation reveals how superconductivity physicist faked blockbuster results](#), by Dan Garisto, Nature, April 2024 .
- [596] [The First Room-Temperature Ambient-Pressure Superconductor](#), by Sukbae Lee, Ji-Hoon Kim, and Young-Wan Kwon, arXiv, July 2023 (22 pages).
- [597] [Superconductor Pb_{10-x}Cu_x\(PO₄\)₆O showing levitation at room temperature and atmospheric pressure and mechanism](#), by Sukbae Lee, Keun Ho Auh, et al., arXiv, August 2023 (18 pages).
- [598] [LK-99: The Live Online Race for a Room-Temperature Superconductor \(Summary\)](#), by Eiri Sanada, 2023.
- [599] [Origin of correlated isolated flat bands in copper-substituted lead phosphate apatite](#), by Sinéad M. Griffin, arXiv, August 2023 (14 pages).
- [600] [The 1/4 Occupied O Atoms Induced Ultraflat Band and the One-Dimensional Channels in the Pb_{10-x}Cu_x\(PO₄\)₆O₄ \(x=0, 0.5\) Crystal](#), by Kun Tao, Chenglong Jia, et al., arXiv, November 2023 (8 pages).
- [601] [Semiconducting transport in Pb_{10-x}Cu_x\(PO₄\)₆O sintered from Pb₂SO₅ and Cu₃P](#), by Li Liu, Zhiqi Liu, et al., arXiv, July 2023 (12 pages).
- [602] [Observation of Zero Resistance Above 100K in Pb_{10-x}Cu_x\(PO₄\)₆O](#), by Qiang Hou, Wei Wei, Xin Zhou, Yue Sun, and Zhixiang Shi, arXiv, August 2023 (7 pages).
- [603] [Ferromagnetic half levitation of LK-99-like synthetic samples](#), by Kaizhen Guo, Yuan Li, and Shuang Jia, arXiv, August 2023 (10 pages).
- [604] [Absence of superconductivity in LK-99 at ambient conditions](#), by Kapil Kumar, N. K. Karn, Yogesh Kumar, and V.P.S. Awana, arXiv, August 2023 (15 pages).
- [605] [First-order transition in Pb_{10-x}Cu_x\(PO₄\)₆O \(0.9 < x < 1.1\) containing Cu₂S](#), by Shilin Zhu, Wei Wu, Zheng Li, and Jianlin Luo, arXiv, August 2023 (7 pages).
- [606] [Academic body invalidates superconductor research results](#), by Shin Ha-Nee and Choi Seo-In, August 2023.
- [607] [Will the LK-99 room temp, ambient pressure superconductivity pre-print replicate before 2025?](#), Manifold, August 2023.
- [608] [Room-temperature superconductor discovered](#), Reddit, 2023.
- [609] [The First Room-Temperature Ambient-Pressure Superconductor](#), comments, PubPeer, July 2023.
- [610] [How hard is it to model LK99 on a quantum computer?](#), by Glenn Jones et al., August 2023.
- [611] [LK-99 isn't a superconductor — how science sleuths solved the mystery](#), by Dan Garisto, Nature, August 2023 .
- [612] [Global Room-Temperature Superconductivity in Graphite](#), by Yakov Kopelevich et al., August 2022 (13 pages).
- [613] [Latest room-temperature claims met with heaps of skepticism](#), by Tim De Chant, TechCrunch, January 2024.
- [614] [Viscosity of Liquid Helium below the \$\lambda\$ -Point](#), by P. Kapitza, Nature, January 1938 .
- [615] [Helium 4](#), by Per Delsing, February 2007 (14 slides).
- [616] [HeLIOS: The Superfluid Helium Ultralight Dark Matter Detector](#), by M. Hirschel, J.P. Davis, et al., arXiv, September 2023 (9 pages).
- [617] [Half-quantum vortices and walls bounded by strings in the polar-distorted phases of topological superfluid ³He](#), by J. T. Mäkinen, V. B. Eltsov, et al., Nature Communications, January 2019.
- [618] [NASA's Cold Atom Laboratory: Four Years of Quantum Science Operations in Space](#), by Kamal Oudrhiri, Robert J. Thompson, et al., arXiv, May 2023 (13 pages).
- [619] [Efimov Physics: a review](#), by Pascal Naidon and Shimpei Endo, arXiv, October 2022 (97 pages).

- [620] Nobel Lecture: Bose-Einstein condensation in a dilute gas, the first 70 years and some recent experiments, by Eric A. Cornell and Carl E. Wieman, NIST, August 2002 (19 pages).
- [621] Heating a dipolar quantum fluid into a solid, by J. Sánchez-Baena, C. Politi, F. Maucher, F. Ferlaino, and T. Pohl, Nature Communications, April 2023.
- [622] The flow of a dense superfluid, by David J. Thouless, 1969 (25 pages) .
- [623] Quantum theory of defects in crystals, by Alexander et al., 1969 (7 pages).
- [624] Probable observation of a supersolid helium phase, by E Kim and M H W Chan, Nature, January 2004 .
- [625] Absence of Supersolidity in Solid Helium in Porous Vycor Glass, by Duk Y. Kim and Moses H. W. Chan, Physical Review Letters, October 2012 (5 pages)  (preprint on arXiv).
- [626] The enigma of supersolidity, by Sébastien Balibar, Nature, 2010 (7 pages).
- [627] Saga of Superfluid Solids, by V.I. Yukalov, arXiv, March 2020 (26 pages).
- [628] Observation of a Dipolar Quantum Gas with Metastable Supersolid Properties, by L. Tanzi, G. Modugno, et al., Physical Review Letters, April 2019 (6 pages)  (preprint on arXiv).
- [629] Long-lived and transient supersolid behaviors in dipolar quantum gases, by L. Chomaz, F. Ferlaino, et al., Physical Review X, March 2019 (17 pages).
- [630] The supersolid phase of matter, by Giovanni Modugno, 2020 (37 slides).
- [631] Theory of the Contribution of Excitons to the Complex Dielectric Constant of Crystals, by Joseph John Hopfield, Physical Review Journal, January 1958 (14 pages) .
- [632] Quantum control of exciton wave functions in 2D semiconductors, by Jenny Hu, March 2024 (8 pages).
- [633] Strongly Nonlinear Nanocavity Exciton Polaritons in Gate-Tunable Monolayer Semiconductors, by Zhi Wang, Physical Review Letters, April 2026 (7 pages).
- [634] Polaritons in van der Waals materials, by D. N. Basov et al., 2016 (9 pages)  (preprint on [escholarship.org](https://arxiv.org/abs/1608.07558)).
- [635] Quantum fluids of light, by Iacopo Carusotto and Cristiano Ciuti, arXiv, October 2012 (76 pages).
- [636] Quantum Fluids of Light, by Iacopo Carusotto, arXiv, November 2022 (8 pages).
- [637] Swimming in a sea of light: the adventure of photon hydrodynamics, by Iacopo Carusotto, 2010 (28 slides).
- [638] Quantum fluids of light, by Jacqueline Bloch, February 2020 (58 slides).
- [639] Quantum Light Fluids - QFL, by ANR, 2016.
- [640] Surface phonon polaritons for infrared optoelectronics, by Christopher R. Gubbin, January 2022 (23 pages).
- [641] Towards polariton blockade of confined exciton-polaritons, by Aymeric Delteil, 2019 (4 pages)  (preprint on arXiv).
- [642] Recent advances in Surface Plasmon Resonance for bio sensing applications and future prospects, by Biplob Mondal et al., August 2020 (31 pages)  (preprint on HAL).
- [643] Quantum enhanced real-time sensing of protein-gold adsorption kinetics, by Mrunal Kamble, Evan Humberd, Tian Li, and Girish S. Agarwal, arXiv, April 2025 (12 pages).
- [644] Polariton: The Krizhanovskii Group, by University of Sheffield, 2022 Accessed via the Wayback Machine, archived June 25, 2022.
- [645] Real-space observation of a two-dimensional skyrmion crystal, by X. Z. Yu, Y. Onose, N. Kanazawa, J. H. Park, J. H. Han, Y. Matsui, N. Nagaosa, and Y. Tokura, Nature, June 2010 (4 pages) .
- [646] Cavity polaritons for new photonic devices, by Esther Wertz, Jacqueline Bloch, Pascale Senellart, et al., Proc. SPIE 7608, Quantum Sensing and Nanophotonic Devices VII, January 2010 (12 pages) .
- [647] Observation of the coupled exciton-photon mode splitting in a semiconductor quantum microcavity, by Claude Weisbuch et al., PRL, July 1992 (4 pages) .
- [648] Observation of three-photon bound states in a quantum nonlinear medium, by Qi-Yu Liang et al., Science, February 2018  (preprint on arXiv).
- [649] Quantum computing with exciton-polariton condensates, by Sanjib Ghosh and Timothy C. H. Liew, npj Quantum Information, February 2020.
- [650] Microcavity Polaritons for Quantum simulation, by Thomas Boulier, Jacqueline Bloch, et al., arXiv, May 2020 (21 pages).
- [651] Polaritonic XY-Ising machine, by Kirill P. Kalinin et al., 2020 (12 pages).
- [652] Enhanced Cavity Optomechanics with Quantum-well Exciton Polaritons, by Nicola Carloni Zambon, Jacqueline Bloch, et al., arXiv, September 2022 (20 pages).
- [653] Acoustic horizons and the Hawking effect in polariton fluids of light, by Elisabeth Giacobino and Maxime J. Jacquet, arXiv, December 2025 (52 pages).
- [654] Floquet Cavity Electromagnonics, by Jing Xu, Changchun Zhong, Xu Han, Dafei Jin, Liang Jiang, and Xufeng Zhang, arXiv, October 2020 (9 pages).
- [655] Dissipation-Based Quantum Sensing of Magnons with a Superconducting Qubit, by S. P. Wolsk, September 2020 (6 pages).
- [656] Quantum Magnonics: Quantum States Generation and Applications, by Zi-Xu Lu, Xuan Zuo, Xin-Lei Hei, Gang Liu, Zeng-Xing Liu, Qi Guo, Peng-Bo Li, and Jie Li, arXiv, July 2026 (42 pages).
- [657] Topological Magnons: A Review, by Paul McClarty, arXiv, June 2021 (12 pages).
- [658] Nonlinear magnon control of atomic spin defects in scalable quantum devices, by Mauricio Bejarano and Helmut Schultheiss et al., arXiv, August 2022 (17 pages).
- [659] Coupling Magnons to an Opto-Electronic Parametric Oscillator, by Junming Wu, Shihao Zhou, Benedetta Flebus, and Wei Zhang, arXiv, October 2025 (9 pages).
- [660] Feedback-enabled magnomechanics in lithium ferrite, by Mehri Ebrahimi, Yunhu Huang, Ali Rashedi, and John P. Davis, arXiv, August 2026 (5 pages).
- [661] The 2020 Skyrmionics Roadmap, by C. Back, Jiadong Zang, et al., arXiv, March 2020 (47 pages).
- [662] Magnon-Skyrmion Hybrid Quantum Systems: Tailoring Interactions via Magnons, by Xue-Feng Pan, Franco Nori, et al., arXiv, April 2024 (9 pages).
- [663] Topological transformation and free-space transport of photonic hopfions, by Yijie Shen, January 2023 (6 pages)  (preprint on arXiv).
- [664] Hopfion rings in a cubic chiral magnet, by Fengshan Zheng, Rafal E. Dunin-Borkowski, et al., Nature, November 2023  (preprint on [dx.doi.org](https://arxiv.org/abs/2311.12345)).
- [665] Space-time crystals from particle-like topological solitons, by Hanqing Zhao and Ivan I. Smalyukh, Nature Materials, September 2025 (20 pages).
- [666] Thermodynamically stable “vortices” in magnetically ordered crystals. The mixed state of magnets, by A. N. Bogdanov et al., 1989 (3 pages).
- [667] Skyrmion Lattice in a Chiral Magnet, by S. Mühlbauer, P. Böni, et al., arXiv, February 2009 (44 pages).
- [668] Skyrmionic SU(6) structured light, by Shinichi Saito, arXiv, March 2025 (18 pages).
- [669] Instantons: thick-wall approximation, by V. F. Mukhanov and A. S. Sorin, arXiv, June 2022 (12 pages).
- [670] Skyrmion-electronics: Writing, deleting, reading and processing magnetic skyrmions toward spintronic applications,

- by Xichao Zhang, Seonghoon Woo, et al., arXiv, November 2019 (80 pages).
- [671] [Mobile Néel skyrmions at room temperature: status and future](#), by Wanjun Jiang, Axel Hoffmann, et al., arXiv, March 2016 (15 pages).
- [672] [Observation of Robust Néel Skyrmions in Metallic PtMnGa](#), by Abhay K. Srivastava, December 2019 (5 pages).
- [673] [Storage and retrieval of optical skyrmions with topological characteristics](#), by Jinwen Wang, Hong Gao, et al., arXiv, December 2025 (7 pages).
- [674] [Skyrmion Logic-In-Memory Architecture for Maximum/Minimum Search](#), by Luca Gnoli, January 2021 (15 pages).
- [675] [Robust and programmable logic-in-memory devices exploiting skyrmion confinement and channeling using local energy barriers](#), by Naveen Sisodia and Olivier Boulle et al., arXiv, May 2022 (11 pages).
- [676] [Neuromorphic weighted sums with magnetic skyrmions](#), by Tristan da Câmara Santa Clara Gomes, Vincent Cros, Julie Grollier, Nicolas Reyren, et al., Nature Electronics, January 2025 (preprint on arXiv).
- [677] [Single skyrmion true random number generator using local dynamics and interaction between skyrmions](#), by Kang Wang, Nature Communications, 2022 (8 pages).
- [678] [The skyrmion switch: turning magnetic skyrmion bubbles on and off with an electric field](#), by Marine Schott et al., 2016 (31 pages) (preprint on arXiv).
- [679] [Magnetic Skyrmion Transistor Gated with Voltage-Controlled Magnetic Anisotropy](#), by Seungmo Yang et al., December 2022 (8 pages).
- [680] [Skyrmion Quantum Diode Prototype: Bridging Micromagnetic Simulations and Quantum Models](#), by Haowen Yang, Yingying Wu, et al., arXiv, January 2026 (18 pages).
- [681] [Skyrmion Qubits: A New Class of Quantum Logic Elements Based on Nanoscale Magnetization](#), by Christina Psaroudaki and Christos Panagopoulos, arXiv, August 2021 (5 pages).
- [682] [Universal quantum computation based on nanoscale skyrmion helicity qubits in frustrated magnets](#), by Jing Xia, Motohiko Ezawa, et al., arXiv, April 2022 (7 pages).
- [683] [Skyrmion-mechanical hybrid quantum systems: Manipulation of skyrmion qubits via phonons](#), by Xue-Feng Pan, Peng-Bo Li, et al., arXiv, April 2024 (16 pages).
- [684] [Nonlinear Field Theory of Large-Spin Heisenberg Antiferromagnets: Semiclassically Quantized Solitons of the One-Dimensional Easy-Axis Neel State](#), by F. D. M. Haldane, Physical Review Letters, April 1983 (4 pages).
- [685] [Topological Quantum Matter](#), by F. Duncan M. Haldane, December 2016 (23 pages).
- [686] [Quantum Spin Hall Effect in Graphene](#), by C.L. Kane and E.J. Mele, arXiv, November 2005 (4 pages).
- [687] [Topological Insulators and the Kane-Mele Invariant: Obstruction and Localisation Theory](#), by Severin Bunk and Richard J. Szabo, arXiv, December 2019 (81 pages).
- [688] [Quantum spin Hall effect: a brief introduction](#), December 2019 (34 slides).
- [689] [Quantum Spin Hall Insulator State in HgTe Quantum Wells](#), by Markus Koenig, Shou-Cheng Zhang, et al., arXiv, October 2007 (16 pages).
- [690] [A topological Dirac insulator in a quantum spin Hall phase \(experimental realization of a 3D Topological Insulator\)](#), by D. Hsieh, M.Z. Hasan, et al., arXiv, October 2009 (12 pages).
- [691] [A complete catalogue of high-quality topological materials](#), by M. G. Vergniory, Zhijun Wang, et al., Nature, February 2019 (6 pages) (preprint on arXiv).
- [692] [All topological bands of all nonmagnetic stoichiometric materials](#), by M. G. Vergniory, Nicolas Regnault, et al., Science, May 2022 (12 pages) (preprint on arXiv).
- [693] [Topological phases of amorphous matter](#), by Adolfo G. Grushin, arXiv, January 2021 (44 pages).
- [694] [Introduction to topological Phases in Condensed Matter](#), by Adolfo G., (28 pages).
- [695] [Detection of Topological Materials with Machine Learning](#), by Nikolas Claussen, B. Andrei Bernevig, and Nicolas Regnault, arXiv, June 2020 (34 pages).
- [696] [Topological Materials : Some Basic Concepts](#), by Ion Garate, 2016 (35 slides).
- [697] [Core Concept: Topological insulators promise computing advances, insights into matter itself](#), by Stephen Ornes, 2016.
- [698] [Topological phases and quasiparticles braiding](#), by Nicholas Read, Physics Today, July 2012 (6 pages).
- [699] [Topological insulator bismuth selenide as a theranostic platform for simultaneous cancer imaging and therapy](#), by Juan Li, Guo-Nan Chen, et al., Scientific Reports, June 2013.
- [700] [Qubit spin ice](#), by Andrew D. King, Alejandro Lopez-Bezanilla, et al., arXiv, July 2021 (18 pages).
- [701] [Topological quantum matter to topological phase conversion: Fundamentals, materials, physical systems for phase conversions, and device applications](#), by Md Mobarak Hosain Polash et al., Materials Science and Engineering: R: Reports, July 2021 (preprint on arXiv).
- [702] [Roadmap on Topological Photonics](#), by Hannah Price, arXiv, 2022 (63 pages).
- [703] [Topological photonic crystals: a review](#), by Hongfei Wang, 2020 (23 pages) (preprint on ncbi.nlm.nih.gov).
- [704] [Topological photonic crystals: physics, designs and applications](#), by Guo-Jing Tang, Jian-Wen Dong, et al., arXiv, January 2022 (60 pages).
- [705] [Topological lasing](#), by Jacqueline Bloch et al., PhLAM Laboratory, Topology and turbulence in fluids of light.
- [706] [Topological-cavity surface-emitting laser](#), by Lechen Yang, Nature Photonics, 2021 (6 pages).
- [707] [Topological insulator vertical-cavity laser array](#), by Alex Dikopoltsev, 2021 (5 pages).
- [708] [Fibonacci Anyons Versus Majorana Fermions: A Monte Carlo Approach to the Compilation of Braid Circuits in \$SU\(2\)_k\$ Anyon Models](#), by Emil G  netay Johansen and Tapio Simula, PRX Quantum, March 2021.
- [709] [Observation of Time-Crystalline Eigenstate Order on a Quantum Processor](#), by Xiao Mi, Matteo Ippoliti, et al., arXiv, August 2021 (16 pages).
- [710] [Realizing topologically ordered states on a quantum processor](#), by K. J. Satzinger, P. Roushan, et al., arXiv, April 2021 (6 pages).
- [711] [Quantum Time Crystals](#), by Frank Wilczek, arXiv, 2012 (6 pages).
- [712] [Classical Time Crystals](#), by Alfred Shapere and Frank Wilczek, arXiv, July 2012 (5 pages).
- [713] [A Decade of Time Crystals: Quo Vadis?](#), by Peter Hanneford and Krzysztof Sacha, arXiv, June 2022 (8 pages).
- [714] [Quantum materials for energy-efficient neuromorphic computing](#), by Axel Hoffmann, Yimei Zhu, et al., arXiv, April 2022 (24 pages).
- [715] [Floquet Time Crystals](#), by Dominic V. Else, Bela Bauer, and Chetan Nayak, Physical Review Letters, August 2016 (preprint on arXiv).
- [716] [Discrete Time Crystals: Rigidity, Criticality, and Realizations](#), by Norman Y. Yao, Andrew C. Potter, Ionut-Dragos Potirniche, and Ashvin Vishwanath, Physical Review Letters, January 2017 (preprint on arXiv).
- [717] [Observation of a continuous time crystal](#), by Phatthamon Kongkhambut, Hans Ke  bler, et al., arXiv, August 2022 (13 pages).
- [718] [Formation of Tesseract Time Crystals on a Quantum Computer](#), by Christopher Sims, arXiv, May 2023 (12 pages).

- [719] [Continuous time crystal coupled to a mechanical mode as a cavity-optomechanics-like platform](#), by J. T. Mäkinen, P. J. Heikkinen, S. Autti, V. V. Zavjalov, and V. B. Eltsov, *Nature Communications*, October 2025 (10 pages).
- [720] [Observation of a Discrete Time Crystal](#), by J. Zhang, Christopher Monroe, et al., arXiv, September 2016 (9 pages).
- [721] [Observation of discrete time-crystalline order in a disordered dipolar many-body system](#), by Soonwon Choi, Mikhail D. Lukin, et al., *Nature*, March 2017 (26 pages).
- [722] [Observation of a many-body-localized discrete time crystal with a programmable spin-based quantum simulator](#), by J. Randall, T. H. Taminiau, et al., arXiv, July 2021 (7 pages).
- [723] [Realization of a discrete time crystal on 57 qubits of a quantum computer](#), by Philipp Frey and Stephan Rachel, arXiv, January 2022 (12 pages).
- [724] [Entanglement boost for extractable work from ensembles of quantum batteries](#), by Robert Alicki and Mark Fannes, *Physical Review E*, April 2013 (4 pages in the related arXiv)  (preprint on arXiv).
- [725] [Colloquium: Quantum batteries](#), by Francesco Campaioli, Stefano Gherardini, James Q. Quach, Marco Polini, and Gian Marcello Andolina, *Reviews of Modern Physics*, July 2024  (preprint on journals.aps.org).
- [726] [Colloquium: Quantum Batteries](#), by Francesco Campaioli, Stefano Gherardini, James Q. Quach, Marco Polini, and Gian Marcello Andolina, arXiv, August 2023 (36 pages).
- [727] [Non-Gaussian enhancement of precision in quantum batteries](#), by Beatriz Polo and Federico Centrone, arXiv, May 2025 (12 pages).
- [728] [The emergence of quantum energy science](#), by Florian Metzler, Jorge I Sandoval, and Nicola Galvanetto, *Journal of Physics: Energy*, October 2023 (11 pages).
- [729] [High-Power Collective Charging of a Solid-State Quantum Battery](#), by Dario Ferraro, *Physical Review Letters*, March 2018  (preprint on arXiv).
- [730] [Subwavelength quantum battery based on an atomic ring array coupled with an incoherently driven atom](#), by Yuxiang Han, Zhengquan Pan, and Hong-rong Li, *Physical Review A*, March 2025 .
- [731] [Superabsorption in an organic microcavity: Toward a quantum battery](#), by James Q. Quach, 2022 (9 pages).
- [732] [Comment on “High-Power Collective Charging of a Solid-State Quantum Battery”](#), by Haowei Xu and Ju Li, arXiv, October 2024 (2 pages).
- [733] [Hyperbolic enhancement of a quantum battery](#), by C. A. Downing and M. S. Ukhtary, arXiv, May 2024 (5 pages).
- [734] [Exceeding the maximum classical energy density in fully charged quantum batteries](#), by Masahiro Hotta and Kazuki Ikeda, arXiv, June 2025 (5 pages).
- [735] [Sizing Up the Potential of Quantum Batteries](#), by Sourav Bhattacharjee, April 2022.
- [736] [Quantum Charging Advantage Cannot Be Extensive Without Global Operations](#), by Ju-Yeon Gyhm, Dominik Šafránek, and Dario Rosa, arXiv, April 2022 (5 pages).
- [737] [Quantum batteries – The future of energy storage?](#), by James Q. Quach, Giulio Cerullo, and Tersilla Virgili, arXiv, October 2023 (9 pages).
- [738] [The battery capacity of energy-storing quantum systems](#), by Xue Yang, Ming-Xing Luo, et al., arXiv, July 2023 (7 pages).
- [739] [Performance of quantum batteries with correlated and uncorrelated chargers](#), by Mohammad B. Arjmandi, Abbas Shokri, Esfandiyar Faizi, and Hamidreza Mohammadi, arXiv, November 2022 (9 pages).
- [740] [Experimental verification of quantum battery capacity with an optical platform](#), by Xue Yang, Ming-Xing Luo, et al., arXiv, December 2024 (15 pages).
- [741] [Controlling Energy Storage Crossing Quantum Phase Transitions in an Integrable Spin Quantum Battery](#), by Riccardo Grazi, *Physical Review Letters*, November 2024  (preprint on arXiv).
- [742] [Charging free fermion quantum batteries](#), by Riccardo Grazi, Fabio Cavaliere, Maura Sassetti, Dario Ferraro, and Niccolò Traverso Ziani, arXiv, March 2025 (32 pages).
- [743] [Strongly correlated perovskite lithium ion shuttles](#), by Yifei Sun, 2018 (6 pages).
- [744] [Micromasers as Quantum Batteries](#), by Vahid Shaghghi, Varinder Singh, Giuliano Benenti, and Dario Rosa, arXiv, April 2022 (6 pages).
- [745] [Toward High-Temperature Light-Induced Spin-State Trapping in Spin-Crossover Materials: The Interplay of Collective and Molecular Effects](#), by M. Nadeem et al., May 2022 (55 pages).
- [746] [Quantum advantage in charging cavity and spin batteries by repeated interactions](#), by Raffaele Salvia, Géraldine Haack, Stefan Nimmrichter, et al., arXiv, April 2022 (4 pages).
- [747] [Remote Charging and Degradation Suppression for the Quantum Battery](#), by Wan-Lu Song, Hai-Bin Liu, Bin Zhou, Wan-Li Yang, and Jun-Hong An, arXiv, February 2024 (13 pages).
- [748] [Enhanced Charging in Multi-Battery Systems by Nonreciprocity](#), by Hua-Wei Zhao, Yong Xie, Xinyao Huang, and Guo-Feng Zhang, arXiv, March 2025 (7 pages).
- [749] [Towards a nuclear isomer quantum battery](#), by Ying-Bo Gao and Fu-Quan Dou, arXiv, May 2026 (5 pages).
- [750] [Charging quantum batteries via Otto machines: The influence of monitoring](#), by Jeongrak Son, Peter Talkner, and Juzar Thingna, arXiv, September 2022 (13 pages).
- [751] [Experimental demonstration of a scalable room-temperature quantum battery](#), by Kieran Hymas, James Q. Quach, et al., arXiv, January 2025 (53 pages).
- [752] [How quantum batteries could lead to EVs that go a million miles between charges](#), by Tristan Greene, *The Next Web*, June 2022.
- [753] [A Wide Range of Testing Results on an Excellent Lithium-Ion Cell Chemistry to be used as Benchmarks for New Battery Technologies](#), by Jessie E. Harlow, 2019 (15 pages).
- [754] [Powering Quantum Computation with Quantum Batteries](#), by Yaniv Kurman, *Physical Review X*, January 2026 (17 pages).
- [755] [Quantum Charging Advantage in Superconducting Solid-State Batteries](#), by Chang-Kang Hu, Dapeng Yu, et al., arXiv, February 2026 (17 pages).
- [756] [IBM quantum platforms: a quantum battery perspective](#), by Giulia Gemme, Michele Grossi, Dario Ferraro, Sofia Vallecorsa, and Maura Sassetti, arXiv, April 2022 (13 pages).
- [757] [Optimal charging of a superconducting quantum battery](#), by Chang-Kang Hu, Dapeng Yu, et al., arXiv, August 2022 (5 pages).
- [758] [Topological quantum materials for energy conversion and storage](#), by Huixia Luo, Peifeng Yu, Guowei Li, and Kai Yan, *Nature Reviews Physics*, July 2022  (preprint on arXiv).
- [759] [Quantum Capacitor: A Coherence-Based Quantum Energy Storage Device](#), by Saeed Haddadi, arXiv, May 2026 (9 pages).

This **Quantum matter** subsection contains 254 bibliographical references and notes containing at least 4,149 pages.

3.11 Extreme quantum

Beyond the basics of quantum physics, many other branches of quantum physics deserve to be examined in this book. They can have various impacts on quantum technologies, noticeably on quantum sensing. They relate to the various attempts to unify all the fields of physics, from quantum physics to relativistic physics and high-energy particle physics. They also deal with explaining the origins and the very structure of the Universe. We will deal with dark matter and dark energy in a later part positioned in the quantum sensing section of this book. Some of these theories are unfortunately used by many false sciences and scams that we will discuss in the section dedicated to quantum pseudosciences and hoaxes.

3.11.1 Quantum field theory

Quantum Field Theory (QFT, given that, later on, we'll use the QFT acronym with another meaning, Quantum Fourier Transform) is a branch of quantum physics that deals with the physics of elementary particles in the relativistic realm, including their creation or disappearance during various interactions, such as electron and positron pairs. These phenomena are generally reproduced in particle accelerators.

QFT also covers the mechanisms of condensed matter such as Bose-Einstein condensates or superfluid helium and more generally, the behavior of quasiparticles, complex collective behaviors such as Cooper's (electron) pairs in superconducting materials.

QFT combines elements of quantum mechanics, special relativity, and classical notions of electromagnetic fields. It is based on a mathematical formalism that is even more difficult to assimilate than the one of non-relativistic quantum physics. It exploits the notion of Lagrangian and Lagrangian integrals over time describing the evolution of fields and the interactions between the fields of several particles.

QFT is used to explain or model several things that are worth telling apart. Spin and the leading fine structure of the hydrogen atom are already accounted for by relativistic quantum mechanics, Dirac's 1928 equation combining the relativistic kinetic correction, spin-orbit coupling and the Darwin term into a splitting of relative order α^2 , meaning about 5×10^{-5} of the gross structure. QED, as the quantum field theory of charged particles and photons, supplies the radiative corrections on top of that, notably the Lamb shift which lifts the $2s_{1/2}$ and $2p_{1/2}$ degeneracy that Dirac's theory predicts and the electron anomalous magnetic moment. QFT also describes the spontaneous emission of photons by atoms during their return to their fundamental state and the interactions behind radioactivity.

The foundations of QFT were created by many scientists starting in 1928: **Paul Dirac**, **Wolfgang Pauli**,

Vladimir Fock (1898-1974, Russian), **Shin'ichirō Tomonaga** (1906-1979, Japanese), **Julian Schwinger** (1918-1994, American), **Richard Feynman** and **Freeman John Dyson** (1923-2020, American).

Dyson gave rise to the notion of the Dyson sphere, which dimensions the level of technological control of energy sources by extraterrestrial civilizations, with a sphere capturing the totality of a star's energy. Shin'ichirō Tomonaga, Julian Schwinger and Richard Feynman received the 1965 Nobel prize in Physics for their work on quantum electrodynamics which is part of QFT.

In the early 1950s, they solved the problem of infinite energy values generated by the initial QFT models by using an adjustment technique called **renormalization**.

Physicists are still struggling to integrate the theory of general relativity into the QFT, preventing it from becoming a "theory of everything" explaining all known physical phenomena in the Universe.

QFT is a theoretical framework, among others, that is applied in four main areas:

- In the physics of **high-energy particles** explored in particle accelerators such as the CERN LHC. It has been supplemented on this point by the standard model that we will see below.
- In the **physics of condensed matter** with superconductivity, superfluidity and the quantum Hall effect, where QFT methods describe collective excitations and quasiparticles.
- In **QED** (quantum electrodynamics), the relativistic quantum field theory of charged particles and photons launched by Paul Dirac in 1928, which studies in particular the production of positrons and positron/electron interactions (attraction, annihilation, pair creation, Compton effect). The **CQED** (cavity QED) sub-branch studies the interactions between atoms and photons confined in optical or microwave cavities. Its solid-state transposition, **cQED** (circuit QED), replaces the atom by a superconducting artificial atom and the cavity by a microwave resonator. This is the framework used by physicists working on superconducting qubits [760].
- In **cosmology** to contribute to modeling the origin and evolution of the Universe as well as certain mechanisms of interaction between black holes and quantum fields.

3.11.2 Quantum vacuum fluctuation

One of the consequences of QFT is the notion of quantum vacuum fluctuation, also called vacuum energy. Building on Heisenberg's indeterminacy principle, QFT models zero-point fluctuations, or vacuum energy, the minimum energy level of quantum systems, with quantum fields fluctuating perpetually even in the absence of any particle.

Q What is the difference between a field and a wave?

In physics, a field is a concept used to describe how a physical quantity is distributed across space and time. It assigns values like temperature, pressure, or force to points in space. It can be scalar, like temperature which has only a magnitude, or vector, like wind velocity which has both a magnitude and a direction. More complex fields, like tensor fields, appear in theories such as general relativity. In quantum field theory, fields are considered fundamental entities, and particles are seen as excitations or disturbances within these fields.

An electromagnetic field is a type of physical field that arises from electric charges which are stationary or in motion. A charge at rest produces an electric field around it. When it moves, it adds a magnetic field. These electric and magnetic fields are intertwined and form the electromagnetic field. It actively influences other charged particles in its proximity. A charged particle within an electromagnetic field will experience a force either pushing or pulling it depending on the field nature and the particle's motion. This interaction is governed by the Lorentz force law, and the behavior of the field itself is described by Maxwell's equations, which are foundational to classical electrodynamics.

A wave, on the other hand, is a dynamic phenomenon, a disturbance or oscillation that travels through space and time, often within a field. Waves carry energy and information, and they manifest in various forms. Mechanical waves like sound or water waves require a medium to travel through. Electromagnetic waves, such as light, propagate through the electromagnetic field and don't need a medium, thus they travel well in vacuum. In quantum mechanics, waves are described by wavefunctions, which represent probabilities rather than physical displacements.

Fields and waves are interrelated. A wave is essentially a pattern or motion within a field. For example, a light wave is a ripple in the electromagnetic field. A field assigns values throughout spacetime and may be static or intrinsically dynamic, its own equations of motion admitting propagating solutions with no external disturbance at all, as with a free electromagnetic wave in vacuum. A wave is a propagating or oscillatory solution of one or several fields. The distinction is therefore between an object and one class of its configurations, not between something static and something dynamic.

In this framework, the ground state of a quantum field is not the classical empty box. Its field amplitude has a nonzero variance and its correlation functions do not vanish, so a total vacuum in the classical sense cannot exist. That state is however stationary and it does not radiate. Real electromagnetic waves are not spontaneously created out of it, otherwise both energy conservation and the observed darkness of empty space would be in trouble.

Under the influence of a surrounding electromagnetic field, this leads to a polarization of the vacuum. The latter even leads to make vacuum birefringent, its refractive index depending on the polarization of the light that gets through it. The phenomenon is however potentially observable only with some very intense electromagnetic fields, approaching the Schwinger critical field of 1.3×10^{18} V/m. The only claimed astrophysical hint so far comes from the optical polarimetry of an isolated neutron star [761].

Theoretical models initially indicated that this vacuum energy would be infinite on the scale of the Universe. They were then corrected using the renormalization method, already mentioned above. Vacuum fluctuations, together with radiation reaction, do account for the spontaneous emission of radiation by the electrons in atoms. Radioactive decay is a different story and each channel has its own interaction. Alpha decay is the quantum tunneling of an alpha cluster through the

combined nuclear and Coulomb barrier, beta decay is mediated by the weak interaction through a virtual W boson, and gamma decay is an electromagnetic transition of the nucleus. In addition to these elementary fluctuations, vacuum is constantly traversed, even in the remotest regions of space, by electromagnetic waves, not to mention the effects of gravitation. The Universe is thus filled with radiations including the cosmological background noise which is a remnant of the big bang, having a temperature of 2.7K. It is the same in a vacuum-packed box because all matter emits radiation.

The concept of vacuum energy originated with **Max Planck** in 1911 when he published an article containing an energy equation for a medium containing a fixed constant, a kind of energy floor for this medium, without being able to interpret it. It was not until 1916 that the chemist **Walther Nernst** (1864-1941, German, who played a key role in launching the Solvay conferences from 1911 onwards) interpreted this constant as the energy level of the vacuum in the absence of any radiation. Its signature shows up when a mode is cooled well below $\hbar\omega/k_B$, about 240 mK for a 5 GHz microwave mode, so that thermal photons vanish and only the residual zero-point contribution remains.

According to the QFT, the Universe is a vast soup containing constantly fluctuating fields, both fermions (leptons and quarks) and bosons (force fields like gluons mediate the strong force that stick together the

quarks that are the elementary constituents of protons and neutrons, and photons, and the cohesion between nucleons is coming from a residual force from strong interactions). This notion of minimum energy level is sometimes compared with the notion of ether, a not completely empty void which dominated 19th century physics, notably for James Clerk Maxwell. The analogy is historical and physically incorrect. The QFT vacuum is Lorentz invariant, which means it has no preferred rest frame and no ordinary viscosity, and a body moving through it at constant velocity feels no drag. That is exactly what separates it from the ether that was ruled out by Michelson and Morley in 1887. Still, these theories are less complete than classical quantum mechanics. One of the solutions is to assume that fermions have a negative vacuum energy and bosons have a positive vacuum energy, both balancing each other. But this has not been demonstrated experimentally, particularly with non-relativistic energy particles.

How do these vacuum quantum fluctuations interact with regular photons? In a Lorentz invariant vacuum, a free on-shell photon does not accumulate stochastic phase, polarization or energy noise along its path. If it did, images of distant quasars propagating over gigaparsec distances would come out blurred and depolarized, which they do not, and that sets extremely tight bounds on any such effect. QED vacuum polarization does modify propagation, but only in strong external fields approaching the Schwinger field, or close to structured boundaries and material interfaces. The optical decoherence measured in real experiments comes from the medium, the components and the detectors, not from the vacuum itself.

Some link could be found between vacuum energy and the dark energy of the Universe as well as gravity [762]. This is very speculative. It could help explain the 68.5% of the energy content of the Universe attributed to dark energy in the Planck 2018 Λ CDM fit [763]. Its density is very low, at about 5×10^{-10} J/m³, or 5×10^{-16} J/cm³.

There are different ways to verify the existence of quantum vacuum fluctuations. The best-known is related to the Casimir effect that we will study in the next part. French and German scientists also managed to interact with this quantum vacuum fluctuation in a semiconductor in 2020 [764].

3.11.3 Casimir effect

The physicist **Hendrik Casimir** (1909-2000, Dutch) predicted in 1948 the existence of an attractive force between two parallel electrically conductive and uncharged plates [765, 766]. He obtained his PhD in 1931 at the University of Leiden in the Netherlands under Paul Ehrenfest. He spent part of his graduate years in Copenhagen with Niels Bohr and was Wolfgang Pauli's assistant at ETH Zurich in 1932-1933, before becoming a professor at Leiden in 1938 and joining Philips Research in 1942. The Casimir effect is interpreted as

being related to the existence of quantum vacuum energy. The experiment imagined by Casimir uses parallel mirrored metal surfaces that are as perfectly flat as possible. They create a Fabry-Perot cavity similar to the one that is used in lasers.




Figure 161: The Casimir effect demonstrated in 1997 between two metallic plates. In the usual picture, the photons created by quantum vacuum fluctuations inside the two mirrors have constrained wavelengths while the photons created outside don't, so the pressure coming from outside is bigger than from inside. This mode-counting picture is heuristic and regulator dependent, the modern derivation going through the field correlation functions and, for real materials, through the Lifshitz formula built on the measured optical response of the plates. Source: [767]. 2023-September 2026.

The Casimir effect is commonly attributed to quantum fluctuations in vacuum. Temporary changes in the energy level at points in the space between the two mirrors would spontaneously generate pairs of very short-lived particles and antiparticles and photons associated with their annihilation. These vacuum fluctuations take place in and out of the volume of the cavity.

Because of the interference effect induced by the cavity, fluctuations at certain frequencies are reduced. The density of electromagnetic energy in the cavity is thus lower than the density of energy outside the cavity as shown in Figure 161. These are spontaneous quantum fluctuations.

The effect cannot be explained by the simple pressure that is higher on the outside than the pressure between the two plates. In detail, the wavelengths of the photons generated by the vacuum outside the plates can be of any size and especially long while inside the plates, only the wavelengths satisfying $\lambda = 2L/n$ fit between the two mirrors, with L the plate separation and n an integer.

The spontaneous electromagnetic spectrum of the vacuum is therefore wider outside the plates than inside, creating a stronger radiation pressure outside than inside, which tends to slightly make the plates move closer together[767–769]. For two parallel and perfectly conducting mirrors of surface A and a distance L between the two mirrors, the force of attraction between the two mirrors follows the formula on the right. In practice, L

is between $0.2\ \mu\text{m}$ and $5\ \mu\text{m}$ and is usually $1\ \mu\text{m}$, which is small and not a macroscopic scale.

$$F_{\text{Cas}} = \frac{\hbar c \pi^2 A}{240 L^4}$$

According to Heisenberg's principle, which is used to explain the effect, energy and time can be linked by the formula on the right. It shows indirectly that over a very short time, a large energy uncertainty is allowed. Strictly speaking this is not one of Heisenberg's canonical relations since time is a parameter and not an operator in quantum mechanics, but it gives the right orders of magnitude.

$$\Delta E \cdot \Delta t \geq \frac{\hbar}{2}$$

Nothing in this enables energy harvesting. The $\Delta E \cdot \Delta t$ relation bounds the spread of energy of a state against the timescale over which an observable changes appreciably. It does not enable a violation of energy conservation over short intervals. Energy is exactly conserved in QED at every order of the perturbative expansion. So, be skeptic when hearing anyone claiming they can harvest energy from vacuum to produce free electricity.

The Casimir force itself is well established and has been measured many times in several geometries. What remains under discussion is the precision of the comparison with theory, at the percent level, which stays sensitive to surface roughness, electrostatic patch potentials, the optical response used in the Lifshitz calculation, geometry corrections beyond the proximity force approximation and thermal corrections, this last point still being debated between the Drude and plasma prescriptions for the low-frequency response of metals. The first experiments validating the Casimir effect were done almost 50 years after the definition of this effect. The first one is that of **Steve Lamoreaux** (American), published in January 1997, using a torsion pendulum with a gold-coated spherical lens of about $12.5\ \text{cm}$ radius facing a flat plate, over separations of $0.6\ \mu\text{m}$ to $6\ \mu\text{m}$ [770]. This sphere-plane geometry sidesteps the parallelism problem that plagues flat plates, at the cost of relying on the proximity force approximation. A real parallel-plate measurement came only in 2002 with Giacomo Bressi et al., and at a better 15% accuracy [771]. The difficulty is about canceling out all the other forces between the two surfaces and they are all much larger than the Casimir effect, particularly electrostatic and van der Waals forces.

His measurement agreed with the prediction at the 5% level. The precision instruments used then detected forces of about one billionth of a Newton. The measurement was improved in 1998 by Umar Mohideen and Anushree Roy with an atomic force microscope and an aluminum-coated polystyrene sphere $196\ \mu\text{m}$ in diameter facing a flat plate, over separations of $0.1\ \mu\text{m}$ to 0.9

μm , bringing the deviation between models and measurements down to 1% at the shortest distances, which remains significant in physics [772]. Gold coatings, micro-mechanical torsion oscillators and nanomembrane setups came afterwards (Figure 164) [773, 774].

The same vacuum fluctuations that produce the Casimir effect also account for other commonly observed physical phenomena such as the electron's anomalous magnetic moment and the Lamb shift, although these are computed directly in QED and are not derived from the Casimir effect itself. The first phenomenon describes a drift of this magnetic moment with respect to Dirac's equation.

The second comes from **Willis Eugene Lamb** (1913-2008, American), Nobel prize in Physics in 1955, who had done his thesis under the supervision of Robert Oppenheimer. Lamb shift is the energy gap observed between the $2s_{1/2}$ and $2p_{1/2}$ levels of the hydrogen atom, which Dirac's theory predicts to be exactly degenerate. The effect is explained by the perturbations coming from vacuum fluctuations and affecting the electron in these two neighboring energy levels, creating the spontaneous generation of photons that are rapidly absorbed by the electron.

The Lamb shift effect was discovered in 1947 by Willis Eugene Lamb and interpreted the same year by **Hans Bethe** (1906-2005, German) for the hydrogen spectrum using the idea of mass renormalization. It was used in the development of post-war quantum electrodynamics.

Vacuum polarization explains only a small part of this shift, contributing about $-27\ \text{MHz}$ out of a total of $1,057\ \text{MHz}$, the bulk coming from the electron self-energy [775]. The calculations are perturbative expansions in the **fine structure constant** α (about $1/137$), the dimensionless coupling constant quantifying the strength of the electromagnetic interaction between elementary charged particles, which also sets the size of the electron's anomalous magnetic moment. It is now known to better than 10^{-10} in relative terms [97, 776].

There is also a **Dynamic Casimir Effect** (DCE), predicted by **Gerald Moore** in 1970, out of his 1969 PhD thesis. It generates real photon pairs out of the vacuum by moving the mirrors used in the Casimir experiment [778, 779]. It was observed for the first time in 2011 by a Chalmers team who replaced the moving mirror by a SQUID modulating the boundary condition of a superconducting transmission line at about $11\ \text{GHz}$, which is equivalent to a mirror moving at a few percent of the speed of light [780].

As with the Casimir Effect, the energy observed is infinitesimal. For the energy to be significant, the mirrors would have to move at relativistic velocities, which is not very practical. And there is no problem with energy conservation, the necessary energy being provided by the mirror movement. The vacuum simply serves as a nonlinear medium!




Figure 162: Explanation of the Lamb shift, an energy shift observed between the $2s_{1/2}$ and $2p_{1/2}$ levels of the hydrogen atom that Dirac's theory predicts as degenerate, explained by quantum vacuum fluctuations impacting electrons. Source: [777].

The interpretation of the Casimir effect is still debated. Some physicists explain it by other mechanisms than vacuum energy. They rely on the **van der Waals** (1837-1923, another Dutch) forces, where atoms attract or repel each other depending on their distances [781–784]. The two are however not rival mechanisms living at different scales. Casimir and van der Waals forces are the retarded and the nonretarded regimes of one and the same electromagnetic dispersion interaction, described continuously by Lifshitz theory, the crossover sitting near the wavelength associated with the dominant absorption of the materials, a few tens of nanometers for metals. Both are measured over nanometer to micrometer separations, so a micrometer gap does not make the Casimir effect macroscopic in any ordinary sense.




Figure 164: Setup of the same experiment, with the gold sphere, the piezo actuator driving the nanomembrane and the fiber interferometer reading its displacement at 1310 nm. Source: [773].

French physicists are quite active in the field, and, in particular, **Astrid Lambrecht**, formerly director of the INP of the CNRS, the Institute of Physics which




Figure 163: Scanning electron microscope view of the gold-coated sphere facing the nanomembrane used for the Casimir force and in situ surface potential measurement. Source: [773].

oversees the physics laboratories of CNRS and now at FZ Jülich and RWTH Aachen in Germany [785] as well as Serge Reynaud (ENS LKB in France). Together in 2026 with other researchers, they described the effect of Casimir zero-point fluctuation in salted water. It could impact our understanding of biology [786].

The Casimir effect could be of interest in quantum metrology to create sensors and in particular NEMS/MEMS, a route already demonstrated in 2001 with a micromechanical torsional actuator driven by the Casimir force [787]. Others are investigating the teleportation of small chunks of energy between qubits but not to the point of implementing real energy harvesting [788–790]. Current and charge could even be teleported [791].

These theories on quantum vacuum fluctuation and the Casimir effect are also fraudulently exploited by the creators of so-called machines capable of capturing vacuum energy, which collect nothing at all in practice. The fluctuation-dissipation theorem ensures that quantum vacuum fluctuations do not violate the second law of thermodynamics. No energy can be recovered thanks to these fluctuations! Forget it. It was the conclusion of a report from the DIA in 2010 [792].

For example, **David Lewis Anderson**, who started the **Anderson Institute** in 1990, claims to use the Casimir effect to travel back in time and create a “free” electricity generator. In other cases, the Casimir effect is exploited in a scientific but borderline way to imagine science fiction scenarios like ways to cross wormholes [793, 794].

NASA even explored the idea of using sails and vacuum fluctuation to propel a space vessel between 1996 and 2002, as one of the ideas of the fancy Breakthrough Propulsion Physics (BPP) program, to no avail. The whole program was awarded a tiny budget of \$1.2M and was later cancelled.

Casimir Inc
USA, 2023, \$14.5M



Casimir Inc was created by Harold White, a former NASA advanced propulsion researcher who spent nearly two decades at the agency, including as director of the EagleWorks Lab where he led experimental work on the EM drive and Alcubierre warp metric concepts (unrelated to the above BPP project). The leadership team includes Jerry Vera (Chief Engineer) and Leonard Dudzinski (SVP Product Development). It was incubated at the Limitless Space Institute, a nonprofit co-founded by Kam Ghaffarian, who is a Casimir investor and board member.

The company's product is the MicroSparc chip, a 5 mm x 5 mm semiconductor device claimed to generate 1.5 V at 25 μ A of continuous electrical power by harvesting energy from quantum vacuum fields via engineered Casimir cavities with metal plates spaced approximately 100 nm apart. Micropillar antennas between the cavity walls are said to collect electrons that quantum-tunnel asymmetrically from the walls, creating a net current without batteries, cords, or charging. Initial target markets are ultra-low-power IoT applications with a stated roadmap to electric vehicles and residential power. Commercial availability is targeted for 2028. The theoretical basis cited by the company is a peer-reviewed paper by White et al. published in March 2026 in *Physical Review Research*, which demonstrates an isospectral acoustic mapping of hydrogen's quantum spectrum onto a dynamic vacuum model. It however contains no experimental energy extraction data [795]. The company came out of stealth in May 2026 with a \$12M seed round led by Scout Ventures. The very premise of the product is highly questionable.

Jovion
USA, 2023



Jovion is another startup willing to extract energy from vacuum. Their principle was developed by Bernard Haisch and Garret Modeld from the University of Colorado Boulder. It uses stochastic electrodynamics (SED) to explain how energy can be extracted from the Zero-Point Energy (ZPE) quantum vacuum using gas cycling and Casimir cavities with hydrogen. They made experiments, subject to huge caution for their interpretation, that could extract power at the picowatt to microwatt range [796]. The authors themselves acknowledge that their interpretation would violate the second law of thermodynamics and that the papers have drawn essentially no response from the community.

While the individual components of the theory such as the Casimir effect and zero-point energy (ZPE) are real, foundational, and measured phenomena in quantum mechanics, the mechanism proposed to extract net usable energy from them is viewed by the scientific community

as a conceptual and thermodynamic impossibility. To extract useful work from any system, there must be a gradient, i.e. a difference between a higher energy state and a lower energy state, or a hot reservoir and a cold reservoir. Because the ZPE vacuum is already at the lowest possible energy state, it cannot act as a fuel source. And Jovion must spend mechanical work pumping the gas through restrictive sub-micron channels, more than the system could extract from orbital transitions.

3.11.4 Theories of everything

The quest for a **theory of everything** has kept many physicists busy for over a century. Its goal would be to consolidate all physics theories in a single formalism and in particular, quantum physics, relativity and gravity. In addition to the QFT, a very large number of explanatory and unifying theories of physics have been developed.

No such theory is considered today as being complete. Figure 165 shows a rough map showing how these different theories are related.

Quantum chromodynamics provides a description of the strong interactions binding quarks together via gluons to form particles called hadrons, which cover the baryons like protons and neutrons and the mesons. Murray Gell-Mann (1929-2019, American, Nobel prize in Physics in 1969) and George Zweig (1937, born in Moscow then American, former PhD student of Richard Feynman, who named them "aces") each proposed the existence of quarks in 1964. Quantum chromodynamics is an extension of quantum field theory formulated in 1972-1973 by Harald Fritzsch, Murray Gell-Mann and Heinrich Leutwyler, and made predictive the same year by the discovery of asymptotic freedom by David Gross, Frank Wilczek and David Politzer, who shared the 2004 Nobel prize [797, 798].

Standard model describes the architecture of known elementary particles and their interactions. It models the electromagnetic, weak and strong fundamental forces. It only lacks gravity to be complete. This model predicted the existence of quarks, these massive particles forming neutrons and protons, in addition to other elementary particles such as the famous Higgs boson whose existence was proven at CERN's LHC in 2012. The expression "standard model" was created in 1975. It relies on a gauge theory because of its mathematical symmetries.

It is not the first of its kind because Maxwell's electromagnetism is also a gauge theory, between magnetic and electric fields. The standard model particles do not cover the famous dark matter whose nature is not yet known. It could be made of particles such as WIMPs (weakly interacting massive particles) or axions, whose mass would sit between 10^{-12} and 10^{-9} of the electron mass, corresponding to the μ eV to meV window ex-

plored by current searches. Various projects have been launched to detect axions using a way to convert them into photons of various energies (X-rays, microwaves, ...) [799].

String theory combines general relativity and quantum physics to propose a quantum explanation of gravity, using a new massless particle, the graviton. According to this theory, elementary particles are tiny strings, open or closed, with vibration types defining the nature of the particle. Their size is of the order of magnitude of 10^{-35} m, the Planck length. According to this theory, the Universe would be a set of vibrating strings.

The graviton would join the three other force carriers: photons for electromagnetism and gluons for the strong interaction that links quarks together in protons and neutrons, both massless, and the W and Z bosons for the weak interaction that governs atomic nuclei and in particular radioactivity, which are massive at 80 and 91 GeV/c² [800, 801]. A proton has two up quarks and one down quark. A neutron has two down quarks and one up quark. Inside a nucleus, a quark can change flavor through the weak interaction, in a two-step chain. In a β^+ -type conversion, an up quark emits a virtual W^+ boson and becomes a down quark, $u \rightarrow d + W^{+*}$, and the W^+ then decays into a positron and an electron neutrino, $W^{+*} \rightarrow e^+ + \nu_e$. In β^- decay, $d \rightarrow u + W^{-*}$ is followed by $W^{-*} \rightarrow e^- + \bar{\nu}_e$. The W is the intermediate mediator of the interaction and never an additional final product. A quark, like an electron, shows no measurable internal structure, with experimental upper bounds on its radius below 10^{-18} m.

Radioactivity comes in three flavors with three distinct mechanisms. Alpha decay emits a particle made of two protons and two neutrons (a helium 4 atom without its electrons) through the quantum tunneling of a preformed alpha cluster across the combined nuclear and Coulomb barrier, a mechanism explained by George Gamow and, independently, by Ronald Gurney and Edward Condon in 1928, and which accounts for the Geiger-Nuttall relation between decay energy and half-life [802]. Beta rays are the electrons or positrons generated by the weak interaction as just described. Gamma rays are photons of very high energy level emitted during an electromagnetic transition of the nucleus. Bosonic string theory, the original version, has only bosonic spacetime excitations, requires 26 dimensions and contains a tachyon that makes its vacuum unstable, which is why it is not viable on its own.

Superstring theory extends string theory by adding worldsheet supersymmetry, which removes the tachyon and brings fermionic states into the spacetime spectrum, something bosonic string theory cannot produce. Neither statement should be read as a particle inventory, one theory covering all observed bosons and the other adding all observed fermions. Both are candidate frameworks and neither has been matched to the particle content of the standard model. It tries to consolidate the description of all forces in a single unified theory. It

quantifies gravity and ties it to other forces. It is based on the notion of supersymmetry which extends the standard model by making each type of boson correspond to a type of fermion. Its distant ancestor is Werner Heisenberg's 1943 S-matrix theory, followed by the dual resonance models of the late 1960s. The supersymmetric versions appeared in 1971 and the field took off in 1984 with the Green-Schwarz anomaly cancellation. It uses 10 dimensions to describe physics, far beyond the four classical dimensions (three for position and one for time). It also uses the notion of "branes", extended objects of various dimensions that generalize the point particle, living in these multidimensional spaces. However, this theory is not unique since there are five variants, which some people try to unify in the **M-theory**, which is based on 11 dimensions. A never-ending story!

Loop quantum gravity theory is another tentative to explain gravity with a quantum model. It discretizes the effects of gravity by presenting space as a meshed structure with quantized areas and volumes of space, and gravitational field quanta connected to each other by links characterized by a spin (that has nothing to do with usual particles spin). For this theory created in the 1980s, the Universe would be a gigantic spin foam. Its main promoters are **Carlo Rovelli** (emeritus at the Center for Theoretical Physics in Marseille and now also at the Perimeter Institute and Western University in Canada [803]) and **Lee Smolin** (Perimeter Institute for Theoretical Physics in Waterloo [804]). The seeds of the theory date back to 1952, with many intermediate stages as shown in Figure 166. It is, above all, a mathematical and topological model. Experimental proposals are starting to show up to verify these theories [805]. It is used to model that the big bang was coming after a big bounce in a cyclical phenomenon with contractions and expansions. It may be possible to detect some fossil signatures of these phenomena. Other theories on gravity avoid the quantization of the gravitational field while keeping consistency with quantum physics and quantum field theory [806].

There are now some proposals to test whether gravity is quantized. One is to use high-precision torsion pendulum at millikelvin temperatures with two tiny diamond crystals, one placed in a quantum superposition, to test gravitational interaction [807]. Another proposal is using an interferometer creating a spatial quantum superposition of a massive object [808]. The reference schemes here are the two 2017 proposals of Sougato Bose et al. and of Chiara Marletto and Vlatko Vedral, where two masses in adjacent matter-wave interferometers become entangled through their mutual gravitational attraction, entanglement being impossible with a purely classical mediator [809, 810].

Sachdev-Ye-Kitaev (SYK) model is a system of N Majorana fermions with random interactions and strongly chaotic dynamics, which at low energy admits a holographically dual description as two-dimensional Jackiw-Teitelboim gravity, a model of two-dimensional "dila-




Figure 165: Approximate classification of quantum physics theories and unification theories. © Olivier Ezratty, 2020-2025.

A brief history of quantum gravity:

- 1952 Flat space quantization (Rosenfeld, Pauli, Fierz, Gupta,...)
- 1959 Canonical structure of general relativity (Dirac, Bergmann, Arnowit, Deser, Misner)
- 1964 Penrose introduces the idea of spin networks
- 1967 Wheeler-DeWitt equation
- 1974 Hawking radiation and black hole entropy
- 1984 String theory
- 1986 New variables for general relativity (Ashtekar, Sen)
- 1988 Loop representation and solutions to the Wheeler-DeWitt equation (Jacobson, Smolin)
- 1989 Extra dimensions from string theory
- 1995 Hilbert space of loop quantum gravity, geometric operators
- 2000' Spin foam models, group field theory, loop quantum cosmology,...

Figure 166: History of quantum gravity theories and models. Source: [811].

ton” gravity used in the context of quantum gravity and holography, named after physicists Roman Jackiw and Claudio Teitelboim, who developed the model independently in the 1980s. The model derives from the 1993 Sachdev-Ye spin liquid, was revisited by Alexei Kitaev in 2015 and formalized by Juan Maldacena and Douglas Stanford in 2016 [812]. Some quantum algorithms have been created to compute this special state of matter and validate quantum gravity theories [813–815]. SYK models are implemented in typical NISQ algorithm use cases on top of Fermi-Hubbard models that are used with condensed matter.

Time travel and tachyons deal with hypothetical particles which would travel faster than light [816].

One overarching trend in high-energy physics driven theories of everything is to consider that everything is made of fields, and massive particles are just artefacts of these fields. Even electrons would be fields [817]. Not far from fields, energy could also be a primitive ontology for the physical world [818].

These are only a few of the many theories being devised. Some freelance researchers also try to create their own “theory of everything”, without usually obtaining any or much feedback from the scientific community [819]. This is the case with the independent hypothesis from Xiaodong Chen [820] and Gunther Kletetschka [821] according to which you need to use three dimensions of time to unify the various fields of physics. Chen does that within the frame of quantum physics to explain

various phenomena like non-locality with one classical infinite time scale and two additional finite time scales. Kletetschka uses three dimensions of time across different scales, the quantum scale (which is responsible for mass generation, operating at the Planck scale), the interaction scale (at the crossroads between quantum and classical phenomena, corresponding to our macroscopic scale) and the large cosmological scale (where gravitational waves sit). In both cases, the additional time dimensions are not seemingly associated with some ontology (common meaning).

Another school of thought considers that, due to non-locality, quantum physics and the theory of relativity are irreconcilable, like with Tim Maudlin. In these two frameworks, time is handled differently. Quantum mechanics treats it as absolute and external, as a fixed parameter while general relativity considers time as

being dynamic and relative, woven into the fabric of spacetime. Tim Maudlin considers that relativity theory must be revised, not quantum mechanics theory [822].

Quantum theories are also created about time moving forward or backward, like the one proposed by researchers from the University of Surrey in the UK in 2025 [823]. This “time-reversal symmetry” is explained here by the Markovian dynamics (stochastic processes with no memory effect) of quantum systems coupled to an environment consisting of a large number of harmonic oscillators. It shows systems that exhibit a time symmetry between their past and future evolution, with two opposite ‘arrows of time’. But this seems to be a mathematical artefact, not a system that actually evolves ontologically back in time. An equivalent theory was built in 2014 at the cosmological scale [824].



3.11.5 Bibliography and notes

These are the bibliographical references and notes for the **Extreme quantum physics** subsection.

- [97] [Determination of the fine-structure constant with an accuracy of 81 parts per trillion](#), by Léo Morel, Zhibin Yao, Pierre Cladé, and Saïda Guellati-Khélifa, *Nature*, December 2020 (5 pages) (preprint on [HAL](#)).
- [760] [Cavity quantum electrodynamics for superconducting electrical circuits: an architecture for quantum computation](#), by Alexandre Blais, Ren-Shou Huang, Andreas Wallraff, Steven M. Girvin, and Robert J. Schoelkopf, *Physical Review A*, February 2004 (14 pages) (preprint on [arXiv](#)).
- [761] [Evidence for vacuum birefringence from the first optical polarimetry measurement of the isolated neutron star RX J1856.5-3754](#), by Roberto P. Mignani, Vincenzo Testa, Daniel González Caniulef, Roberto Taverna, Roberto Turolla, Silvia Zane, and Kinwah Wu, *Monthly Notices of the Royal Astronomical Society*, October 2016 (9 pages) (preprint on [arXiv](#)).
- [762] [Casimir cosmology](#), by Ulf Leonhardt, *arXiv*, February 2022 (41 pages).
- [763] [Planck 2018 results. VI. Cosmological parameters](#), by Nabila Aghanim, Yashar Akrami, Mark Ashdown, Jonathan Aumont, and Carlo Baccigalupi, *Astronomy and Astrophysics*, September 2020 (67 pages).
- [764] [Electric field correlation measurements on the electromagnetic vacuum state](#), by Ileana-Cristina Benea-Chelms, Francesca Fabiana Settembrini, Giacomo Scalari, and Jérôme Faist, *arXiv*, September 2020 (17 pages).
- [765] [On the attraction between two perfectly conducting plates](#), by Hendrik Casimir, 1948 (3 pages).
- [766] [Electromagnetic vacuum fluctuations, Casimir and Van der Waals forces](#), by Cyriaque Genet, 2004 (18 pages).
- [767] [The Casimir Effect](#), by Kyle Kingsbury, 2014 (82 slides).
- [768] [The Casimir effect and the physical vacuum](#), by G. Takács, 2014 (111 slides).
- [769] [Zero-Point Energy and Casimir Effect](#), by Gerold Gründler, 2013 (47 pages).
- [770] [Demonstration of the Casimir Force in the 0.6 to 6 micrometers Range](#), by Steve K. Lamoreaux, *Physical Review Letters*, January 1997 (4 pages) .
- [771] [Measurement of the Casimir Force between Parallel Metallic Surfaces](#), by Giacomo Bressi, Giovanni Carugno, Roberto Onofrio, and Giuseppe Ruoso, *Physical Review Letters*, March 2002 (4 pages) (preprint on [arXiv](#)).
- [772] [Precision Measurement of the Casimir Force from 0.1 to 0.9 micrometers](#), by Umar Mohideen and Anushree Roy, *Physical Review Letters*, November 1998 (4 pages) (preprint on [arXiv](#)).
- [773] [Casimir Force and In Situ Surface Potential Measurements on Nanomembranes](#), by Steve Lamoreaux et al., 2012 (6 pages) (preprint on [arXiv](#)).
- [774] [The Casimir force between real materials: experiment and theory](#), by Galina L. Klimchitskaya, Umar Mohideen, and Vladimir M. Mostepanenko, *Reviews of Modern Physics*, February 2009 (59 pages).
- [775] [The Vacuum Polarisation Contribution to the Lamb Shift Using Non-Relativistic Quantum Electrodynamics](#), by Jonas Frafjord, 2016 (61 pages).
- [776] [Measurement of the Electron Magnetic Moment](#), by Xing Fan, T. G. Myers, B. A. D. Sukra, and Gerald Gabrielse, *Physical Review Letters*, February 2023 (6 pages) .
- [777] [The Lamb Shift](#), by Carl R. Nave, *HyperPhysics*, Department of Physics and Astronomy, Georgia State University, September 2026.
- [778] [Quantum Theory of the Electromagnetic Field in a Variable-Length One-Dimensional Cavity](#), by Gerald T. Moore, *Journal of Mathematical Physics*, September 1970 (13 pages) .
- [779] [Electro-mechanical Casimir effect](#), by Mikel Sanz, 2018 (10 pages).
- [780] [Observation of the Dynamical Casimir Effect in a Superconducting Circuit](#), by Christopher M. Wilson, Göran Johansson, Arsalan Pourkabirian, Michaël Simoen, J. Robert Johansson, Tim Duty, Franco Nori, and Per Delsing, *Nature*, November 2011 (12 pages) (preprint on [arXiv](#)).
- [781] [The origin of Casimir effect: Vacuum energy or van der Waals force?](#), by Hrvoje Nikolic, 2018 (41 slides).
- [782] [The Casimir-Effect: No Manifestation of Zero-Point Energy](#), by Gerold Gründler, *arXiv*, August 2013 (15 pages).
- [783] [All wrong with the Casimir effect](#), by Astrid Karnassnigg, 2014 (3 pages).
- [784] [The Casimir effect: a force from nothing](#), by Astrid Lambrecht, 2007 (5 pages).

- [785] [The Casimir effect theories and experiments](#), by Romain Guérout et al., 2010 (28 slides).
- [786] [Casimir effect with dielectric matter in salted water and implications at the cell scale](#), by Larissa Inácio, Felipe S. S. Rosa, Astrid Lambrecht, Paulo A. Maia Neto, and Serge Reynaud, arXiv, January 2026 (13 pages).
- [787] [Quantum Mechanical Actuation of Microelectromechanical Systems by the Casimir Force](#), by Ho Bun Chan, Vladimir A. Aksyuk, Raymond N. Kleiman, David J. Bishop, and Federico Capasso, Science, March 2001 (4 pages) .
- [788] [The Quest to Use Quantum Mechanics to Pull Energy Out of Nothing](#), by Charlie Wood, May 2023.
- [789] [A review of applications of Quantum Energy Teleportation: from experimental tests to thermodynamics and space-time engineering](#), by Boris Ragula and Eduardo Martín-Martínez, arXiv, July 2025 (27 pages).
- [790] [Quantum Energy Teleportation without Limit of Distance](#), by Masahiro Hotta, Jiro Matsumoto, and Go Yusa, arXiv, January 2014 (7 pages).
- [791] [Beyond Energy: Teleporting Current, Charge, and More](#), by Kazuki Ikeda, arXiv, July 2024 (17 pages).
- [792] [Concepts for extracting energy from the quantum vacuum](#), by DIA, Defense Intelligence Agency, 2010 (58 pages).
- [793] [One Theory Beyond the Standard Model Could Allow Wormholes that You Could Actually Fly Through - Universe Today](#), by Matt Williams, August 2020.
- [794] [Humanly traversable wormholes](#), by Juan Maldacena and Alexey Milekhin, arXiv, November 2020 (3 pages).
- [795] [Emergent quantization from a dynamic vacuum](#), by Harold White, Physical Review Research, March 2026 (7 pages).
- [796] [Casimir-cavity-induced conductance changes](#), by Garret Model, Ayendra Weerakkody, David Doroski, and Dylan Bartusiak, Physical Review Research, April 2021 (7 pages).
- [797] [50 Years of Quantum Chromodynamics](#), by Franz Gross, Xiaorong Zhou, et al., arXiv, December 2022 (567 pages).
- [798] [Advantages of the Color Octet Gluon Picture](#), by Harald Fritzsch, Murray Gell-Mann, and Heinrich Leutwyler, Physics Letters B, November 1973 (4 pages) .
- [799] [Dark Matter](#), by Marco Cirelli, Alessandro Strumia, and Jure Zupan, arXiv, July 2024 (515 pages).
- [800] [Graviton detection and the quantization of gravity](#), by Daniel Carney, Valerie Domcke, and Nicholas L. Rodd, arXiv, August 2023 (7 pages).
- [801] [Quantum decoherence of gravitational waves](#), by Hiroki Takeda and Takahiro Tanaka, arXiv, June 2025 (23 pages).
- [802] [Zur Quantentheorie des Atomkernes](#), by George Gamow, Zeitschrift für Physik, March 1928 .
- [803] [Princeton seminars on physics and philosophy](#), by Carlo Rovelli, arXiv, July 2024 (51 pages).
- [804] [Lee Smolin Public Lecture Special: Einstein's Unfinished Revolution](#), by Lee Smolin, 2019.
- [805] [Testing the Quantumness of Gravity without Entanglement](#), by Ludovico Lami et al., PRX, May 2024 (47 pages).
- [806] [A Postquantum Theory of Classical Gravity?](#), by Jonathan Oppenheim, PRX, December 2023 (37 pages).
- [807] [Probing the Quantum Nature of Gravity through Classical Diffusion](#), by Oliviero Angeli, Angelo Bassi, et al., arXiv, March 2025 (17 pages).
- [808] [Testing Whether Gravity Acts as a Quantum Entity When Measured](#), by Farhan Hanif et al., Physical Review Letters, October 2024.
- [809] [Spin Entanglement Witness for Quantum Gravity](#), by Sougato Bose, Anupam Mazumdar, Gavin W. Morley, Hendrik Ulbricht, Marko Toroš, Mauro Paternostro, Andrew Geraci, Peter Barker, Myungshik S. Kim, and Gerard Milburn, Physical Review Letters, July 2017.
- [810] [Gravitationally-induced entanglement between two massive particles is sufficient evidence of quantum effects in gravity](#), by Chiara Marletto and Vlatko Vedral, Physical Review Letters, December 2017  (preprint on [arXiv](#)).
- [811] [The philosophy behind loop quantum gravity](#), by Marc Geiller, 2001 (65 slides).
- [812] [Remarks on the Sachdev-Ye-Kitaev model](#), by Juan Maldacena and Douglas Stanford, Physical Review D, November 2016  (preprint on [arXiv](#)).
- [813] [Quantum chaos in the sparse SYK model](#), by Patrick Orman, Hrant Gharibyan, and John Preskill, arXiv, September 2024 (22 pages).
- [814] [Sachdev-Ye-Kitaev model on a noisy quantum computer](#), by Muhammad Asaduzzaman, Raghav G. Jha, and Bharath Sambasivam, arXiv, May 2024 (6 pages).
- [815] [Simulating sparse SYK model with a randomized algorithm on a trapped-ion quantum computer](#), by Etienne Granet, Yuta Kikuchi, Henrik Dreyer, and Enrico Rinaldi, arXiv, July 2025 (9 pages).
- [816] [Testing Tachyon-Dominated Cosmology with Type Ia Supernovae](#), by Samuel H. Kramer and Ian H. Redmount, arXiv, March 2024 (20 pages).
- [817] [Schrödinger's Field Equation](#), by Jacek Mroczkowski, arXiv, May 2025 (8 pages).
- [818] [Energy as a Primitive Ontology for the Physical World](#), by J.E. Horvath and B.B. Martins, arXiv, June 2025 (14 pages).
- [819] [Unified Theory Research Team](#), by Pierre and Frédéric Lepeltier, September 2021.
- [820] [3-Dimensional time: the physics behind quantum mechanics and unified interactions](#), by Xiaodong Chen, arXiv, October 2005 (47 pages).
- [821] [Three-Dimensional Time: A Mathematical Framework for Fundamental Physics](#), by Gunther Kletetschka, World Scientific Connect, April 2025 (14 pages).
- [822] [The Great Rift in Physics](#), by Tim Maudlin, arXiv, March 2025 (28 pages).
- [823] [Emergence of opposing arrows of time in open quantum systems](#), by Thomas Guff, Andrea Rocco, et al., Scientific Reports, January 2025 (18 pages).
- [824] [Identification of a Gravitational Arrow of Time](#), by Julian Barbour, Physical Review Letters, October 2014 (6 pages)  (preprint on [arXiv](#)).

This **Extreme quantum physics** subsection contains 66 bibliographical references and notes containing at least 2,026 pages.

3.12 Quantum foundations

Philosophy is a process of critical reflection and questioning about the world, knowledge and human existence. It creates a connection between all these dimensions. The discovery of quantum physics at the beginning of the 20th century created a real philosophical shock wave through the upheavals it brought to our understanding of the world at the microscopic level.

In practice, these upheavals occur mainly at the nanoscopic scale, that of atoms and their constituents, the nuclei and electrons. However, quantum effects can also be observed at the scale of large groups of particles that can be mesoscopic or even macroscopic, as is the case with large molecules and their wave-particle duality, in Bose-Einstein condensates or superconducting currents. Knowing in all this that the frontier between quantum physics and classical physics has regularly evolved over the last century.

It called into question key notions such as the links between “reality” (what is) and observations (“what we can see of it”), or between ontology and epistemology (how do we know what we know?). And the debate is still raging about it. If you meet a group of quantum physicists and want to have some fun, ask them simple questions like “*what is a quantum state?*”, “*what does the wavefunction mean?*” or “*what is contextuality?*”! They may not agree on a (simple) answer.

3.12.1 Quantum physics and its missing ontology

Science has always been closely linked to philosophy. It is not by chance that a doctorate is a “PhD”, or Doctor of Philosophy, whether in humanities or in so-called hard or exact sciences. The great physicists and mathematicians of the 19th and early 20th centuries were also philosophers, which is less common now, due to specialization.

The creators of quantum physics were constantly questioning the impact and meaning of their discoveries. **Niels Bohr** was also both a physicist and a philosopher, influenced in particular by **Søren Kierkegaard** (1813-1855, Danish). **Erwin Schrödinger** was even more of a philosopher than a physicist. He had studied Western and Indian philosophy before creating the famous wave equation that bears his name. Michel Bitbol indicates that in the epilogue of “*What is life? Mind and matter*”, Erwin Schrödinger wondered whether consciousness was singular or plural. If consciousness is only experienced in the singular, its extension to a global consciousness such as that of the Universe is only a risky extrapolation and difficult to prove experimentally. The thesis of a consciousness of the Universe is defended by some scientists [825, 826].

An assistant to Niels Bohr, **Werner Heisenberg** also invested a lot of his time in philosophy and it related

well with his work around the mathematical modeling of quantum physics and the famous indeterminacy principle.

The debates that agitated quantum physics scientists often took as much the form of philosophical jousting as of physical or mathematical debates, even more since the founders of quantum physics were usually not experimenters and were rather theoreticians (Figure 167).

Popular history has, moreover, forgotten the names of the experimentalists. We mentioned many of them at the beginning of this book, such as Johann Balmer, Theodore Lyman, Friedrich Paschen, James Chadwick, Arthur Holly Compton, George Paget Thomson, Clinton Davisson and Lester Germer. The names of these experimental physicists generally do not ring a bell to the general public and scientists, whereas the general public has heard much more about Max Planck (with his constant more than for the black body radiation explanation), Albert Einstein (for the theory of relativity more than for the photoelectric effect explanation), Niels Bohr (for his model of the atom), Erwin Schrödinger (more for his cat analogy than for his wave equation) and Werner Heisenberg (for his indeterminacy principle, commonly called uncertainty, but not much for this huge work on quantum physics mathematical foundations). Among the founding fathers, Paul Dirac, Wolfgang Pauli and John von Neumann were geniuses but are way behind in notoriety.

Quantum physics generated endless debates since its beginnings because its formalism is difficult to associate with the principles of reality usually applicable in classical physics. Intuitive classical physics understanding has historically been associated with an ontology.

In Newtonian physics, the notion of state with position and motion of an object and the laws of evolution of these properties allow the prediction of phenomena such as the motion of planets. These evolutions are perfectly observable and deterministic.

Quantum physics was created without such an ontology although it explained some known physical phenomena like the blackbody radiation, the photoelectric effect or hydrogen’s absorption and emission spectrum. An ontology deals with what is, types and structures of objects, properties, events, processes and relationships in all areas of reality. It is usually opposed to or complementary to epistemology, which deals with how we obtain valid knowledge, mostly through observations.

Quantum physics was formulated mainly as a set of mathematical postulates that could help predict experimental results. You mainly have the Schrödinger wavefunction for non-relativistic massive particles and others like Dirac and Klein-Gordon equations for relativistic particles, and then the second-quantization formulation for quantum photonics and for high-energy particle physics. In quantum physics, the prediction instrument is a probabilistic wave function that is difficult to apprehend. It is coupled with a whole host of

new notions that have no equivalent in the macroscopic and classical world: energy quantization, wave-particle duality which applies to matter (electrons, atoms) and photons (all have a momentum p related to a wavevector k using the reduced Planck constant, as $p = \hbar k$), the influence of measurement on the quantities to be measured, measurement indeterminacy and the notion of chance [827].

Quantum physics is a predictive, not descriptive theory. It does not describe physically the electrons and other particles when they behave quantumly. It does not physically explain entanglement nor wave-particle duality, but it does relate them to experiments and calculations.

Einstein's position was that quantum physics was an incomplete theory when creating his famous EPR paradox in 1935. Werner Heisenberg asserted in 1927 that quantum physics established the final failure of causality! The knowledge of the present no longer made it possible to predict with certainty the future from the application of the laws of physics, even more so as the knowledge of the present with precision may be impossible, though this is also a debated question. *"In the strong formulation of the causal law, 'If we know the present with exactitude, we can predict the future, it is not the conclusion, but rather the premise that is false. We cannot know, as a matter of principle, the present in all its details'"* [828].

Some like Niels Bohr concluded that it was useless and even counterproductive to create some quantum physics ontology. Many attempts were contradicting each other or weren't even realist per se. Some like Hugh Everett believed that reality was defined by a universal $|\Psi\rangle$ function, David Bohm devised some pilot waves explanations, Christopher Fuchs et al. are focused on the role of agents actions and experiences in quantum Bayesianism (QBism), the more recent approach called 'Contexts, Systems and Modalities' (CSM) argues that physical objects are systems within contexts, and so on. We end up having competing postulated ontologies frequently enabling the same predictions [829–831]. These are hard to sort out. Debates abound between "psi-ontic" and "psi-epistemic" views of quantum physics, psi being related to the mathematical description of a quantum system state [829, 832–835]. Engineers and realists usually prefer an ontic-based interpretation that tries to describe the quantum world independently from the observer - but what and where is the real physical object [836, 837]?

The relationship between measured values, measurement and the observer is also discussed. Would a true measurement be one that does not alter the quantity to be measured at all, a feat hard to attain in the infinitely small? It is often said that quantum mechanics is contextual, the measurement depends on its context, which does not detract from its objectivity. This approach is challenged by the Bayesian quantum interpretation (QBism for Quantum Bayesianism) promoted from 2002

onwards by Carlton Caves, Christopher Fuchs, Rüdiger Schack and then David Mermin [838, 839].

The mathematical formalism of quantum physics from 1900 to 1935 was not at all disconnected from the observable physical world. It made it possible to explain experimentally studied phenomena such as black body radiation, interference from Young's slits with light and matter waves, from spectral excitation lines of atoms under a wide range of conditions to solid state properties.

We have seen how important they are in photonics, with trapped ions, cold atoms or NV centers. Electron spin retrospectively explained the two spots observed in 1922 in the Stern-Gerlach experiment [840], which were then interpreted as the space quantization of the orbital angular momentum, electron spin being proposed only in 1925. This experiment made a beam of heated silver atoms pass through a non-homogeneous magnetic field, which generated two distinct spots on a screen. Hyperfine levels have another origin, the coupling between the electron and the nuclear spins.

Relativistic quantum chemistry derived from Paul Dirac's equations explained spectral shifts of transitions involving inner-shell electrons of heavy atoms moving at relativistic velocities. The list is long.

If quantum physics explained experimental measurements, linking the observed reality and the theory, it was however insufficient to produce an unanimously accepted representation of reality. It is part of a history of science that described matter step by step, with nested Russian dolls. Atoms were initially abstract, theoretical entities before being embodied and accurately described and then directly observed as we do now with electron microscopes or cryogenic microscopy (Cryo-EM), or by trapping and observing individual ions or atoms.

The very existence of atoms was not agreed upon at the end of the 19th century between Ludwig Boltzmann who believed in them and Wilhelm Ostwald and Ernst Mach who opposed them. Protons and neutrons were then discovered. These were split into quarks and gluons with particle accelerators, turning the physical world into a maybe endless fractal, though there are reasons to think the opposite. Obstacles to understanding it could simply be related to the enormous amount of energy that needs to be injected into particle accelerators, which is increasing the more elementary the particles are.

3.12.2 Quantum foundations questions

Quantum physics philosophy belongs to the broad field of **quantum foundations**. It focuses essentially on the multiple possible interpretations of the same theory and its mathematical formalism.

The related questions are multiple, such as:




Figure 167: Various books on quantum physics and philosophy. Updated in 2024.

- **Nature of reality.** The Copenhagen interpretation considers the wave function to represent potentialities [841], while Many-Worlds sees all possibilities as real in separate branches.
 - **Role of the observer.** Does reality exist independently of the observer? Psi-ontic interpretations think so, while psi-epistemic interpretations favor the role of the observer (Figure 168). See more on Wigner’s friend thought experiment.
 - **Determinism and indeterminism** are about elucidating the measurement problem and its indeterminism and the associated notions of causality [488]. Contextuality theorems exclude broad classes of ontological models in which operationally equivalent preparations or measurements are assigned context-independent underlying representations. They constrain the kind of pre-existing values a hidden-variable model may assign, under their stated assumptions, rather than establishing on their own that the measured information literally did not exist beforehand.
 - **Wave function ontology** is about describing the physical meaning of the wave function. Is it a real wave of an indeterminate nature or is it a simple statistical and probabilistic mathematical tool, incomplete in its ability to describe physical reality? These different interpretations can be evaluated using the criteria of scientificity from Karl Popper (1902-1994, Austrian/English), according to which a theory is scientific if it can be refuted by crucial experiments giving precise results. The theory cannot be shown to be irrefutable. A proven scientific theory is therefore always caught between two worlds, as a theory corroborated by experimental facts, until proven otherwise. The history of physics has shown, however, that the “serious” theories of the past were mainly challenged by the broadening of their perspective and context: with large masses and high velocities (for relativity) and in the microscopic (for quantum physics). In their initial contexts, they remained perfectly valid. I like the very current example of the search for dark matter, which would represent about 85% of the matter in the Universe and about 27% of its total energy density. Its existence is not yet experimentally demonstrated but is assumed by the application of the laws of gravity and relativity applied to the cohesion of galaxies. It can be refuted or partially verified at present in at least three ways: by discovering elementary particles associated with dark matter (quantum detectors are built in this sense, and have so far not generated any signals), by modifying the laws of general relativity as the Israeli Mordehai Milgrom is trying to do with MOND [842], or by discovering hidden matter such as the dust of galaxies that could explain all or part of their cohesion without using dark matter. Beliefs in God, some supernatural force, and many areas of metaphysics do not belong to science because they are neither demonstrable nor refutable.
 - **Realism.** Scientific realism is the view that the world described by science exists independently of our observations or perceptions, well-confirmed scientific theories are approximately true, even about unobservable entities like atoms, electrons, photons, black holes, or quarks and the success of science and predictive and explanatory power is best explained by assuming that its theories genuinely describe reality. In the case of quantum physics, the debate is centered around the effects of quantum measurement. Classical realism assumes that quantum objects like electrons have definite properties whether or not we observe them, or at least, they are defined by their probabilistic wave function or quantum state [843]. According to quantum physics, measurement affects the system with the wave function collapse. So, is the wavefunction or the system quantum state (in a given measurement basis) a description of reality or just a mathematical tool used to predict a measurement outcome (in the given basis)? Is there a reality while we only apply unitary transformation (quantum gates in quantum computing) to our quantum objects before measurement-induced collapse [844]?
- In relation to the wave function ontology, how could we define the “reality” of a maximally entangled state during a quantum computing circuit execution and its evolution over time, way before we run some quantum measurement at the end of the circuit, and not

considering the ancilla qubit measurement done for quantum error correction? Is this some “inaccessible reality”? Is the reality of the physical state of the quantum computer defined only by its measurement, yielding only bits of information, much less information than what the state vector representing the qubit register contains? Is the “reality” only about the measurement result which is a pale shadow of the wealth of the register content? Can we have some reality without an observer? Can we isolate the physicality of the quantum system from the information it is holding while running some quantum circuit? This deals with understanding what the quantum system wavefunction represents, just some information or a quantum physical state with an ontological existence in space and time, aka an “ontic” state. And we’re back to the first item of this list.

Gerard 't Hooft is a theoretician in high-energy physics who was awarded the Nobel prize in physics in 1999. He is saying that quantum physics should be revisited beyond its statistical framework and not taken for granted as it is today [845]. He even argues that state superpositions are not real. In a word, he is the extreme case of a “realist”. It seems to be in line with the proposal from Vasil Yordanov to consider that the motion of electrons in atoms follow some Brownian or stochastic motion [846], a position not that far from Jacob Barandes’s work [847].

- **Nonlocality** is about explaining entanglement, local realism and the completeness or incompleteness of quantum mechanics [848].

Let’s dig into some details...

Wigner’s friend paradox is a thought experiment related to the second question, proposed by Eugene Wigner in 1961 [849–851]. It extends the Schrödinger’s cat paradox thought experiment. It explores the nature of observation and the role of the observer’s so-called consciousness in the measurement process. In the scenario, a laboratory contains a friend (Wigner’s friend) who is observing a quantum system, such as a photon that can be in a superposition of two states (e.g., horizontal or vertical polarization). The friend measures the photon’s polarization, which according to quantum mechanics, causes the wave function to collapse to a definite state. An observer outside the lab, or a “super-observer”, does not know the result of the measurement. From his perspective, the entire laboratory, including his friend and the photon, is said to be in a superposition state and both are entangled. And there are many variations of this thought experiment [852].

This has some connections with the notion of **intersubjectivity** which refers to the mutual recognition and validation of experiences between “conscious” beings. It is about how we align our perceptions of reality. This theme is highly present in discussions about QBism [853, 854]. Intersubjectivity’s role in the Wigner friend experiment is the following: if quantum states collapse only

upon observation, whose observation counts? Wigner’s? His friend’s? Both? What is the shared reality between the observers? Intersubjectivity demands that observers agree on what happened. The paradox exposes a tension between the subjective experience of measurement and the objective description of quantum systems. The problem here is: agree on what? The device settings? The polarizer angle? The software configuration of the measurement devices? The noise and imperfections from the whole measurement chain? It is not really explained [855].

Some quantum physicists like Renato Renner (ETH Zurich) are toying with the paradox, up to associating the result of the measurement to the consciousness of the observer [856, 857], leading to the Frauchiger-Renner (FR) paradox [858]. What the paper shows is that a specified set of assumptions cannot all hold at once: quantum theory applies universally, including to agents who themselves use it, agents may adopt each other’s conclusions consistently, and each measurement yields a single definite outcome. Running their protocol, agents reach contradictory conclusions about the same run. Interpretations disagree on which of these assumptions, or which step of the agents’ reasoning, has to go.

In 2023, an experiment was even proposed by Howard Wiseman to validate a no-go theorem that creates a contradiction between four metaphysical assumptions (local agency, physical supervenience, ego absolutism, and friendliness) and replacing Wigner’s friend by a quantum computer running a sort of human level artificial intelligence [859, 860]. I would advise them to replace this complicated highly futuristic setting by a simple Raspberry Pi driven system. Also, about the so-called observer consciousness impact on measurement, let’s play with having 100 observers looking at a screen presenting a measurement result. Would they see the same result? If so, which is obvious, which observer’s consciousness is influencing the measurement and its results? Some researchers are even connecting Wigner experiments with black hole physics and their evaporation [861].

This notion of the influence of consciousness really lacks any scientific ground. It doesn’t define consciousness from a physical standpoint and leads to very esoteric debates [862, 863]. How is it interacting with the experiment? How does consciousness interact with the wavefunction of the measured quantum system? What physical laws govern this? And it’s even not falsifiable. You cannot build an experiment to test whether consciousness collapses the wavefunction. There is no evidence at all. On top of this, all these observers and their nested boxes are classical and not quantum. They can’t be defined by a simple wave function.

In a 2026 proposal, Maria Violaris asserted that quantum observers could communicate across multiverse branches [864]. This thought experiment turned into a real experiment with 5 IBM qubits which was not con-

clusive one way or the other to validate one quantum physics interpretation [865].

Local Friendliness no-go theorem challenges the assumptions about observers, locality, and realism, particularly in the context of extended Wigner’s friend (EWF) paradox scenarios [866, 867]. Derived in 2020 by Kok-Wei Bong, Eric Cavalcanti, Howard Wiseman *et al.* [860] and revisited in 2023 with Eleanor Rieffel [859], the theorem combines Bell-type experiments testing nonlocality with Wigner’s friend scenarios where observers are themselves measured. In the 2020 version, Local Friendliness (LF) is the conjunction of absoluteness of observed events (an outcome observed by anyone is a single objective fact, and not a fact relative to an observer or to a branch), locality and no-superdeterminism, these last two being frequently merged into a single local agency assumption stating that a freely chosen setting can only be correlated with events in its future light cone. In that specific formulation, the LF assumptions are strictly weaker than the local hidden variable assumptions of Bell’s theorem, in the precise sense that local hidden variables imply LF but not the converse, so every LF violation is also a Bell violation while the reverse does not hold. The 2023 version replaces the set with local agency, physical “supervenience”, “ego absolutism and friendliness”, and the comparison with Bell has to be restated for that formulation. The authors derive Local Friendliness inequalities which are constraints on correlations that should hold if these assumptions are true. Quantum theory predicts violations of these inequalities, meaning that at least one assumption must be false. This was even tested on an IBM NISQ quantum processor in 2026 [868]. Nothing in the theorem says that an observer’s thoughts or consciousness physically influence a measurement outcome. The friend is a measuring agent, and a photodetector or a quantum computer plays that role just as well. It generates endless debates going way beyond my pay grade [869–871]. It is even connected to the notion of black holes paradoxes [872, 873].

Contextuality is a notion frequently used in quantum physics according to which quantum measurements do not reveal pre-existing values but depend on the measurement context like the angle of a spin measurement in the famous Stern-Gerlach experiment or of a photon polarization angle. Works on contextuality frequently deal with discrete variable quantum systems. It can be extended to continuous variable systems [874, 875].

Several no-go and go theorems are related to contextuality:

- **Kochen-Specker theorem** *aka* the Bell-Kochen-Specker theorem demonstrated by John S. Bell in 1966 [876] and by Simon B. Kochen and Ernst Specker in 1967 [877, 878] is a “no go theorem” that creates constraints on the types of putative hidden-variable theories trying to explain the predictions of quantum mechanics in a context-independent

way [879]. It was demonstrated elegantly by Adán Cabello in 1996 with 18 vectors grouped in 9 orthonormal bases in a dimension four space, meaning two qubits [878, 880–884], and in 2025 in dimension three with a much more symmetrical set requiring only 14 bases, against 16 for the previous Peres and Penrose record [885]. Quantum contextuality requires a Hilbert space of dimension at least 3, so a single qubit is not enough while a qutrit or a pair of qubits are [886].

- **Gleason’s theorem** (1957) is a “go theorem” from which Born’s rule and the projection aspect of quantum measurements can be deduced, for Hilbert spaces of dimension at least three, which is why the single qubit case has to be treated separately [887–889]. The previously quoted Kochen-Specker’s theorem [890] appears as corollary (or consequence) of Gleason’s theorem.
- **Robert Spekkens’s theorem** (2004) is another no-go theorem which extends noncontextuality beyond sharp measurements, to preparations, transformations and unsharp measurements, and shows that quantum theory admits no generalized-noncontextual ontological model [891, 892].

Several contextuality papers revisit and extend its conceptual framework, with [893] which relates generalized contextuality to antidistinguishability, [894] which shows that contextuality can be certified from a set of individually noncontextual experiments, and [895] which analyzes the logical structure linking contextuality and nonclassicality. Karl Svozil recasts contextuality proofs as graph coloring problems [896] while Partha Ghose and Sudip Patra formalize Bohr’s complementarity with a many-valued contextual logic [897]. Other works evaluate the computational complexity of deciding contextuality [898], build a resource theory of contextuality using simplicial methods [899], and show that all entangled states can be certified by a violation of noncontextuality inequalities [900]. Contextuality is used to rule out nonlinear modifications of quantum theory [901] and can be simulated with invasive-measurement maps [902]. On the experimental side, it has been demonstrated on NISQ superconducting processors [903], with cycle scenarios based on Pauli observables providing the most economical test settings [904]. Contextuality is also positioned as a computing resource, being what supplies the “magic” enabling universal quantum computing with magic state distillation [905], and it is involved in the design of quantum error correction codes. Quantum contextuality seems to be an intrinsic characteristic of fault-tolerant quantum codes and protocols, on top of entanglement and magic [906].

Retrodiction is the process of making probabilistic statements about past events based on current information [907]. It is the time-reversed counterpart of prediction. It can help reinterpret quantum states not just as forward-evolving entities, but as carriers of in-

formation about both past and future. While quantum mechanics is formally time-symmetric, measurement introduces an apparent asymmetry. Retrodiction helps explore this asymmetry and whether it's fundamental or epistemic, due to lack of information. Existing retrodiction techniques are related to Bayesian interpretation of probability and use quantum state smoothing [908–910]. It can even detect past entanglement of the examined quantum system [911].

Real or complex numbers. This is another debate which can be considered as being part of quantum foundations. The question is about whether complex numbers are needed for supporting quantum physics and its postulates? Complex number operations can always be computationally translated into real arithmetic [912]. A prominent 2021 theorem by Marc-Olivier Renou, Nicolas Gisin et al. showed that under the standard tensor-product rule for composing independent systems, a strictly real-valued quantum theory fails to reproduce correct outcomes in specific multi-party network experiments, a prediction since confirmed experimentally [913]. Recent research argues that this failure hinges on that composition convention rather than on real numbers as such. By adopting alternative rules for combining systems, a purely real theory could replicate all standard quantum predictions [914, 915]. If the independence of physical sources is defined strictly by what can be observed in a laboratory rather than by rigid mathematical states, real and complex quantum theories remain experimentally indistinguishable [916]. This operational approach requires careful verification against other physical constraints [917]. The debate is therefore about which composition postulate to adopt. It does not establish that complex numbers are merely a convenient bookkeeping device, since under the conventional postulates the two theories are not equivalent and the real formulations recover equivalence only by paying for it elsewhere [918].

3.12.3 Quantum foundations interpretations

Several interpretations of quantum physics have emerged to try to provide answers to these many questions [919–921] (Table 5). Four very different things fall under the word interpretation. An **interpretation of standard quantum theory** makes no new prediction and is empirically equivalent to the textbook formalism, as with Copenhagen, many-worlds, QBism or relational quantum mechanics. A **modified and testable dynamics** changes the equations and can therefore be experimentally constrained or excluded, as with the GRW and CSL collapse models or gravitationally induced collapse. A **speculative research program** is an incomplete framework that makes predictions but has not produced a decisive one yet. An **unsupported claim** is not falsifiable at all and sits outside science, which is where consciousness-collapse and quantum medicine belong. Placing an interpretation in one of these four categories and stating its empirical status comes before evaluat-

ing it. They also sometimes ambition to reconstruct the usual postulates of quantum mechanics [922]. The debates are still open [923, 924]! A July 2025 survey published by Nature did show that out of 1,101 respondents, from 15,582 physicists contacted, meaning a 7% response rate, 36% of the physicists favored the Copenhagen interpretation, 17% information-based approaches, 15% the many worlds interpretation and 7% the Bohm/De Broglie pilot wave interpretation, while 16% wrote in their own answer, either rejecting all the proposed options or considering that no interpretation is needed [925]. An earlier and much smaller poll of 33 participants at a 2011 conference had already produced a comparable spread [926].

These interpretations are described here by historical order of appearance (Table 5).

Quantum Logic was proposed by **John von Neumann** and **Garrett Birkhoff** in 1936 [117]. It emphasizes the logical structure of quantum propositions and measurements, inspired by how classical logic is used in classical physics. Here, propositions correspond to projection operators on a Hilbert space, forming a lattice that captures the logical relationships between propositions.

This framework highlights differences from classical logic with measurement non-commutativity, superposition and entanglement. In quantum logic, a proposition negation corresponds to the orthogonal complement in the Hilbert space. It incorporates conditional probabilities. It is also used in quantum information science, where it helps develop protocols for quantum communication, cryptography, and error correction [927]. Quantum set theory is an extension of quantum logic [928].

Copenhagen interpretation is the most widely taught interpretation of standard quantum theory and the historical default, several other interpretations being equally compatible with the formalism [929]. Its naming was created by Werner Heisenberg in 1955. It consolidated Heisenberg's and others contemporary views on the realism of quantum physics and concepts that were not promoted by Niels Bohr in the 1920s like the wave packet collapse and the views on the subjective observer in quantum measurement. It's a post-mortem consolidation on the quantum physics foundations. It is essentially probabilistic. Quantum physics postulates and the wave function describe all that we can know about reality but not reality itself, which is neither accessible nor meaningful. It adopts the positivist approach according to which one sticks to observations, laws and phenomena, without seeking to know their intrinsic nature.

It was the “Bohrian” side of the historical debate between **Niels Bohr** and Albert Einstein, mainly between 1927 (in the famous Solvay Conference in Brussels) and 1935 (the EPR paradox paper and subsequent discussions). Adopted by Werner Heisenberg, Max Born, Wolfgang Pauli and Paul Dirac, it is the classical and

dominant interpretation of quantum physics that is still mostly taught like in the Cohen-Tannoudji/Diu/Lal   quantum physics bible. It is satisfied with an essentially mathematical and probabilistic model that does not seek to physically describe the entire real world. There are, moreover, sub-branches in the Copenhagen interpretation, notably around the open and closed theories that had opposed Heisenberg and Dirac from 1929 onwards [930].

Bohm interpretation came from **David Bohm** (1917-1992, American then Brazilian and British). He proposed in 1952 a so-called deterministic version of quantum mechanics, labeled “De Broglie-Bohm theory”. It was inspired by ideas initially promoted, but then abandoned, by Louis de Broglie and took up the idea of the existence of hidden variables, insinuated by Albert Einstein in the 1930s, in the form of pilot waves, as part of the so-called “*double solution*”. It was built on the idea that quantum particles have definite positions, their velocity being fixed by a guidance equation derived from the wave [931–933]. In other words, a wave (a field) and a particle (with a position) are interacting together, both potentially playing the role of a hidden variable. This could explain the double-slit experiment interferences but also the Stern Gerlach experiment as well as, later, Bell inequality violation experiments [934]. In that later case, however, there are potential inconsistencies with special relativity due to the existence of non-local hidden variables and Bohmian trajectories in this interpretation although they are really hidden [935]. Also, there’s a debate on the meaning of the wave function in this theory: is it representing a concrete physical entity, or is it “nomological”, like a law of nature, a question sharpened by the Ψ -ontology theorems proved in the framework PBR (Pusey, Barrett and Rudolph) in 2012 [831, 936–938].

The Bohm-De Broglie pilot-wave theory could be applied in various fields of physics that are beyond the scope of this book, like high-energy particle physics, quantum field theory, gravitation, and cosmology [939].

Bohmian and De Broglie’s interpretations have been promoted by physicists like Detlef D  rr (LMU M  nchen, 1951-2021) [940], Nino Zangh   (Universit   di Genova) [941], Shelly (Sheldon) Goldstein (Rutgers University), John W. M. Bush (MIT) [942], Ward Struyve (KU Leuven) [943], Jean Bricmont (UCLouvain, Belgium) [944], Pascal Naidon (RIKEN) [945], Xavier Oriols (Universitat Aut  noma de Barcelona) [946, 947], Angel S. Sanz from Universidad Complutense de Madrid [948], and Aur  lien Drezet (Institut N  el, France), on the heels of some form of superdeterminism driven by time-symmetry [949–953].

Many Worlds Interpretation (MWI) and its Universe wave function. It started in 1957 with **Hugh Everett** (1930-1982, American) proposing the “relative state formulation” of quantum mechanics which made a distinction between absolute and relative states [954]. After being almost forgotten, it was revived by

Bryce DeWitt in 1970 to become the multiple worlds or multiverse interpretation, in an article published in *Physics Today* [123], although a “multiverse” can have a different meaning in cosmology, with several bubbles of “universes” that are so distant that they can’t be seen from each other and therefore, exist independently in parallel but are entirely different [955].

DeWitt’s interpretation is also called EWG for Everett-Wheeler-Graham. John Wheeler was supervisor of Hugh Everett’s thesis and Neill Graham, a student of DeWitt [124, 956, 957]. It is said to be realistic in the sense that the Universe would be a huge wave function with an immensely large number of parameters. It never collapses and the world is deterministic but splits into infinitely many parallel branches, the measurement being described as a quantum branching phenomenon where the universe “branches” into multiple versions, each one representing a different possible outcome (for the gazillion measurements happening everywhere...) [958]. DeWitt’s interpretation transforms quantum probabilities within this universal wave function into parallel worlds that exist simultaneously. Since it is impossible to verify that parallel worlds exist, the theory is not refutable, but counter arguments abound [959]. It is also not obvious to provide an idea of the number of such universes. It would probably not be countable, be enormous and hard to associate with some sound epistemology [960].

We are therefore far from an experimentally supportable interpretation [961]. This theory has also been promoted by David Deutsch, also known for his quantum algorithms and keeps quantum foundation physicists busy. It has a “Everything is a Quantum Wave” interpretation (EQWI) variation [962]. It feeds many science fiction dreams and mad mysticism, everything being linked to everything and vice versa, especially souls and consciousness. Howard Wiseman extended this theory with embedding interactions between these many worlds, *ala* Fringe [963]. The MWI quantum interpretation is quite lively with a lot of promoters and abundant paper production [964–969]. It also has its detractors, like Emily Adlam and Jacob Barandes [970, 971]. Its most complete modern defense remains David Wallace’s book [972].

Twin Worlds is a variant of MWI with only two worlds proposed by Daniel Braun (Germany) in 2026 [973].

Modal interpretation, a class of interpretations created in 1972 by **Bas van Fraassen** which introduced a distinction between a dynamical state over time and a value state at a given time [974].

Consistent histories aka quantum history which avoids the use of a wavefunction collapse to describe physical processes and tries to get rid of the so-called measurement problem [975]. It was proposed by **Robert Griffiths** in 1984.

Dynamic histories which reinterpret quantum mechanics as deterministically evolving dynamical world




Figure 168: Ontology, epistemology, methodology and methods defined. The various grey boxes gather theories or interpretations that are only very loosely related. © Olivier Ezratty, 2024-2025.

lines in a 5D universe not far from a many-worlds interpretation [976, 977]. It emerged in the early 1980s.

GRW theory published in 1986 by the Ghirardi-Rimini-Weber trio proposes a different formulation of Schrödinger’s equation with a spontaneous reduction of the wave function that is not simply related to the notion of measurement [978]. This is one of the rare formalisms that can be experimentally falsified, and it is actively being constrained, with underground experiments having already ruled out the simplest parameter-free version of the gravity-related Diósi-Penrose collapse model [979, 980].

Superdeterminism which suggests that the apparent randomness and indeterminism observed in quantum measurements, particularly with entangled quantum objects, are actually the result of underlying deterministic processes that are influenced by hidden variables, maybe at the deep Planck-energy level, which is not reachable experimentally. These variables are determined in such a way that they ensure the outcomes of quantum experiments and are pre-determined [981–983]. The theory was proposed by **Carl H. Brans** in 1988, the term itself having been coined by John Bell in 1985. Among others, it is promoted by **Sabine Hossenfelder** and **Tim Palmer**, who explained why superdeterministic is not preventing free will to exist [984]. Other researchers argue that it’s more a metaphysical than a scientific proposition [985].

Relational Quantum Mechanics (RQM) is a relational ontology crafted by **Carlo Rovelli** in 1994 and published in 1996 [986], that considers a quantum state to be defined by the relation between any pairs of systems. One of them can be an observer. It creates the notion of relational properties and ‘relative facts’ that

may become approximately ‘stable facts’ if one system is large enough. In this interpretation, physical properties are not intrinsic to quantum objects but are relational. They depend on the interactions between systems. It is inspired by special relativity principles and its observer reference model [987, 988]. This relational aspect is different than what we discussed in 3.2.2, with the electron quantum numbers being properties of electrons in the context of the atoms which host them. Here, the relationship is not just about the electron in its physical context, but also in relationship with the observer. This interpretation spurred some debates with Bohmians like Aurélien Drezet, with [989] responding to [990], responding to his critique [991] which was reacting to [992].

QBism is a derivative from quantum Bayesianism which interprets quantum physics through the eye of the observer agent actions and experiences [993–996]. It started with some ideas by **Edwin Jaynes** (1922-1998, American, yes, the Jaynes from the Jaynes-Cummings Hamiltonian) and promoted by **Christopher Fuchs** [997, 998], **David Mermin** [999] et al., in the early 2000s, based on **Frank Ramsey**’s anti-realist interpretation of probability (1903-1930, British philosopher) and **Ludwig Wittgenstein**’s work (1889-1951, Austrian philosopher) [1000].

QBism emphasizes the subjective nature of quantum states and probabilities and considers that quantum states are not objective descriptions of physical systems but reflect an individual’s personal degree of belief about potential measurement outcomes, based on available information. It adopts a maximalist Bayesian probability model, viewing probabilities as expressions of subjective degrees of belief that are updated according to Bayes’ theorem with new information.

interpretation	year	author(s) and key references	deterministic?	ontic wave function?	unique history?	hidden variables?	collapsing wave function?	observer role?	local dynamics?	counter-factually definite?	extant universal wave function?
ensemble interpretation	1926	Max Born [1001, 1002]	agnostic	no	yes	agnostic	no	no	no	no	no
Copenhagen interpretation	1927	Niels Bohr, Werner Heisenberg [929, 930]	no	some	yes	no	some	no	yes	no	no
de Broglie-Bohm pilot wave theory	1927-1952	Louis de Broglie, David Bohm [931, 933]	yes	yes	yes	yes	phenomenological	no	no	yes	yes
quantum logic	1936	Garrett Birkhoff, John von Neumann [117, 927]	agnostic	agnostic	yes	no	no	interpretational	agnostic	no	no
time-symmetric theories	1955	Satosi Watanabe [907, 1003]	yes	no	yes	yes	no	no	no	no	yes
many-worlds interpretation	1957	Hugh Everett [123, 967, 1004]	yes	yes	no	no	no	no	yes	ill-posed	yes
consciousness causes collapse	1961-1993	John von Neumann, Eugene Wigner, Henry Stapp [1005]	no	yes	yes	no	yes	causal	no	no	yes
many-minds interpretation	1970	H. Dieter Zeh [148]	yes	yes	no	no	no	interpretational	yes	ill-posed	yes
consistent histories	1984	Robert B. Griffiths [975]	no	no	no	no	no	no	yes	no	yes
transactional interpretation	1986	John G. Cramer [1006]	no	yes	yes	no	yes	no	no	yes	no
objective-collapse theories	1986-1989	Ghirardi-Rimini-Weber (GRW) [978], Roger Penrose [1007]	no	yes	yes	no	yes	no	no	no	no
relational interpretation	1994	Carlo Rovelli [986, 987]	no	no	agnostic	no	yes	intrinsic	possibly	no	no
QBism	2010	Christopher Fuchs, Rüdiger Schack [997, 1008]	no	no	agnostic	no	yes	intrinsic	yes	no	no
CSM	2015	Alexia Auffèves, Philippe Grangier [171, 391, 494, 888, 1009–1015]	no	ill-posed	yes	no	updating	context	yes	no	no

Table 5: Fourteen interpretations of quantum physics sorted by order of appearance and compared along nine discriminating criteria. Reading the table by column rather than by row segments the field into five families, using the color code in the first column: statistical and operational readings that refuse an ontic wave function (gray), deterministic hidden-variable and time-symmetric theories (blue), Everettian branching where a universal wave function never collapses (purple), theories where collapse is a physical process (orange), and agent- or context-relative interpretations where a state encodes information rather than being (green). Only the objective-collapse theories of the orange family are in principle falsifiable, since they alone predict departures from the Schrödinger equation, the transactional interpretation sharing the standard quantum predictions. Green cells mean the criterion holds, red that it does not, cream that the interpretation is agnostic or that the question is ill-posed within it. Source: [Interpretations of quantum mechanics](#), Wikipedia, plus Olivier Ezratty additions for CSM and for the references, 2024-2026.

Fundamental entities are observers with measurements being interactions between agents and systems, resulting in personal experiences. It suggests that there is no single, universal quantum state. Each agent has its own quantum state based on its knowledge and experiences. Measurement outcomes are not pre-existing properties but are created during the measurement act, personal to the agent. It rejects the notion of an objective reality independent of observers.

QBism faces criticism for its subjectivity and the challenge of explaining inter-subjective agreement when agents communicate their measurement results, compare and correlate their data.

Category theory applied to quantum physics is frequently associated with **Gennaro Auletta**'s encyclopedic synthesis of quantum foundations [1016], but the categorical quantum mechanics program proper was launched in 2004 by **Samson Abramsky** and **Bob Coecke**, who axiomatized quantum protocols in compact closed categories, an approach that later led to the creation of the ZX-calculus [1017].

Conceptuality Interpretation, formerly Quantum Conceptual Turn starting in 2002, from **Diederik Aerts** and **Massimiliano Sassoli de Bianchi** [1018, 1019]. It considers that quantum entities are non-spatial and non-objectual, and more concepts than physical objects. Measurement is viewed as a contextual interaction, similar to how concepts acquire meaning when placed in cognitive contexts. Superposition, entanglement, and nonlocality are interpreted as manifestations of conceptual relationships. In this interpretation, the quantum world is not made of things, but of meaning-bearing entities that actualize properties through interaction, much like how concepts become concrete in human thought. It is related to the field of quantum cognition which uses quantum formalism to model human decision-making and conceptual processing [1020, 1021].

Topos Quantum Theory was developed starting in 1997 by Chris Isham and later developed with Andreas Döring and Jeremy Butterfield. Using quite a cryptic philosophical vocabulary, it reimagines quantum theory using “topos theory”, a branch of category theory that generalizes set theory and logic. The aim is to create a neo-realist framework for quantum physics. It is based on “daseinisation”, a process that approximates quantum propositions within classical contexts, allowing for a contextual assignment of truth values, then “spectral presheaf” which replaces the traditional state space and encodes all possible measurement contexts, contextual-ity with quantum properties being described relative to commutative subalgebras (*aka* contexts, not far from CSM), not globally. In this framework, quantum systems are described not by wavefunctions in Hilbert space, but by presheaves of sets over a category of contexts [1022–1024].

Quantum Darwinism ambitions to explain how the classical world emerges from the quantum world [1025–

1032]. It was developed by **Wojciech Zurek** et al. between 2003 and 2009, in the continuation of his work on decoherence and einselection [1033]. In this theory, quantum systems interact with their environment, resulting in the proliferation of certain stable states, known as pointer states, which can be redundantly recorded across the environment. These pointer states are robust against decoherence and thus become the “classical” states observed at the macroscopic level. The environment acts as a communication channel, disseminating information about the quantum system to multiple observers without significantly disturbing the system itself. Here, decoherence is described as the “branching of quantum worlds” [1034]. This redundant spreading of information enables a consensus on the state of the system, leading to an objective reality that aligns with classical physics principles. In 2025, a team from Zhejiang University, the University of Houston and Los Alamos illustrated this interpretation with an experimental validation using a processor with 12 superconducting qubits (2 as the studied system, 10 acting as the environment surrounding the system) [1035]. But it didn’t invalidate other interpretations!

CSM was proposed in 2015 by **Alexia Auffèves** and **Philippe Grangier** in order to reconcile the Copenhagen interpretation with realistic models, based on the earlier work from Philippe Grangier on contextual objectivity [171, 391, 888, 1009–1015]. It is a minimalist ontology in which measured properties called **modalities**, are attributed to a **system** (studied system, as isolated as possible) within a **context** (completely specified measurement device and settings). Modalities are jointly associated to the system and its context, not just the system, building a contextual objectivity. Within the usual quantum formalism, a modality is a pure quantum state and a context is implementing a complete set of commuting observables (CSCO). For a given context, CSM defines N distinguishable modalities that are mutually exclusive. If one modality is “true”, or “realized”, the others are “wrong” (or “false”), or “not realized”. The value of N , called the dimension, is a characteristic property of a given quantum system, and is the same regardless of the context (Figure 169). In CSM, randomness is a direct consequence of the quantization postulate and the contextual nature of reality. Otherwise, after a quantum measurement, you are in a basis state of the given context (or eigenstate) and the measurement can be repeated by generating the same result. However, this can only work with non-destructive measurements which is easier with spins measurement “*ala* Stern Gerlach” than with photon polarization measurement using polarizers and photon detectors [1036].

CSM can help explain the origin of probabilities, nonlocality and quantum-classical boundary. CSM also uses the notions of extracontextuality and extravalence which clearly define how context changes and introduce quantum randomness in most cases, and some determinism in some cases.




Figure 169: CSM's simple view. © Olivier Ezratty, 2024.

CSM's ongoing work involves understanding how quantum measurement operates and where the Heisenberg cut (which is the separation between the quantum system being measured from the classical system doing the measurement) occurs with combining Koopman's approach and type III von Neumann algebra with Infinite Tensor Products (ITP) [494]. The Koopmanian approach is a classical systems representation using Hilbert spaces, similar to quantum mechanics, coming from the work of Bernard Koopman in 1931, who showed that classical dynamics can be described using unitary evolution in a Hilbert space [1037–1039]. In that representation *aka* Koopman algebra, classical observables are treated as commutative multiplication operators and time evolutions are governed by a unitary Koopman operator. This allows classical systems to be embedded in a quantum-like framework, enabling hybrid quantum-classical models. It is done within the framework of the Koopman–von Neumann (KvN) theory [491, 1040, 1041], which can be used to create quantum algorithms solving nonlinear dynamics [1042]. Type-III algebras are a class of von Neumann algebra coming from the Murray and von Neumann classification of factors published between 1936 and 1943, that are infinite-dimensional and non-factorizable, lack a trace (as a result of their infiniteness), and commonly arise in quantum field theory and at thermodynamic limits. They can describe systems with infinite degrees of freedom, such as quantum fields or statistical ensembles of quantum objects [1043]. All this work is part of CSM2 while the work until 2019 belongs to CSM1.

Geometrodynamic Model of Reality from Shlomo Barak in 2017.

Quantum Coherentism from Claudio Calosi and Matteo Morganti in 2018 [1044].

Nomic toy theory is another deterministic interpretation of quantum mechanics [1045]. It was introduced in 2024 by Johannes Fankhauser, Tomáš Gonda, and Gemma De les Coves.

Generative Quantum Theory (GQT) by Francisco Pipa in 2024 [1046].

RQD (Relational Quantum Dynamics) proposed by Arash E. Zaghi in 2024 [1047], “*an interpretation that places relationships and information, not isolated ob-*

jects, at the foundation of physics”. The proposal is explicitly positioned against other interpretations.

Constructor Theory was proposed by David Deutsch in 2013 [1048] and later extended by Chiara Marletto [1049], both from the University of Oxford, with additional contributions from Vlatko Vedral, Bob Coecke and Maria Violaris [1050]. It shifts focus from traditional dynamical laws to counterfactual statements, which are descriptions of what transformations are possible or impossible under physical laws, when the circumstances are different but still grounded in known laws and patterns of nature [1051]. Rather than describing how systems evolve over time, Constructor Theory defines tasks as transformations between states and introduces “constructors” entities capable of performing these tasks repeatedly without degradation. This meta-theoretic framework ambitions to unify diverse domains like thermodynamics, quantum computation, and information theory. In the landscape of quantum interpretations, Constructor Theory doesn't directly compete with traditional views like the Copenhagen interpretation, Many-Worlds Interpretation (MWI), or Bohmian mechanics. Instead, it treats quantum mechanics as a subsidiary theory within a broader framework that includes “super-information”, a generalization of quantum information. Constructor Theory is more meta-theoretical, offering a new way to express physical laws that could potentially reconcile quantum mechanics with other domains, including computation and biology. They derived this work for understanding gravity [1052–1054], quantum interference [1055], time [1056] and trying to reconcile quantum mechanics and general relativity [1057]. Some experimental protocols were proposed in 2026 by Marletto, Vedral and Deutsch [1058, 1059].

BHSI (Branched Hilbert Subspace Interpretation) is a proposal from Xing M. Wang of an alternative view of quantum measurement, conceptualizing it as a unitary branching of the local Hilbert space into decoherent, independent, and unitarily evolving subspaces [1060, 1061]. It is positioned relatively to other interpretations. It avoids the wave function collapse from the Copenhagen interpretation and explains the Born rule through branch weights associated with the initial system state. It sidesteps the notion of parallel worlds by entangling branches with the local environment within a single, unified world. In BHSI, the observer's state is updated through causally engaging and disengaging operators, meaning there is only one observer in a single world who experiences one outcome per observation. This interpretation suggests that information is preserved, and measurement-induced decoherence is localized and potentially reversible, allowing for recoherence under controlled conditions. These features, including locally controlled decoherence-recoherence processes (CDRP), could be observable in quantum teleportation and modern Stern-Gerlach style experiments.

Algorithmic Idealism is a theoretical framework proposed by Markus Müller (Vienna) [1062] and Krzysztof

Sienicki (Poland) [1063–1065] that “*reconceptualizes reality as an emergent phenomenon arising from informational and computational processes*”. It “*reinterprets quantum mechanics as a computational process governed by algorithmic probability informational simplicity and utility optimization*”. In this interpretation, “*reality is modeled as an agent environment feedback loop where agents maximize utility by predicting selfstate transitions*”. It reminds me of the theory of everything from Stephen Wolfram which is built on graph structures.

Some other more recent interpretations are **Quantum Reconstruction** by Philipp Berghofer from the University of Graz which extends QBism [1066–1068], **Super Quantum Mechanics** coming from Russia which considers that states in Hilbert space are subject to multiple quadratic constraints, proposed in January 2025 [1069] and the Tlapan Interpretation of Quantum Mechanics (TIQM) proposed by Alejandro Frank from Mexico also in 2025 which reframes quantum collapse as an emergent, thermodynamic phase transition governed by (highly contentious) retrocausal boundary conditions, with testable predictions that bridge quantum mechanics with statistical physics [1070].

Other ontologies abound like **Structuralism**, **Fundamentalism**, **Perspectival Objectivity** [1071], **Pluralism** (atomism), **Monism** and **Infinetism**, but their scope goes beyond quantum physics.

Where are the schools of quantum foundations? In Europe, there is a quantum foundations epicenter in Italy but you find contributors in most countries. There are also two foundations who provide research grants on quantum foundations: the **Templeton Foundation** and **FQXi** (Foundational Questions Institute) which covers quantum foundations and cosmology and was created in 2005 by cosmologists Max Tegmark and Anthony Aguirre.

3.12.4 Other interpretations

Among the physicists who have contributed to the field of quantum physics philosophy, **Pascual Jordan** (1902–1980, German) built a theory of free will according to which one is not freer by acting randomly or in a determined way, breaking the idea that quantum non-determinism would be a proof of human free will.

Henry Stapp (1928, American) worked on consciousness and believes that it governs the world and reality and that it can only be explained by quantum physics [1005]. This is the kind of theory that makes non-falsifiable hypotheses that then become gospel for the quantum medicine quacks. The basic idea is that brain chemistry, like all chemistry, is based on many facets of quantum physics. This becomes complicated when it deals with entanglement to explain consciousness, like what Roger Penrose is doing with the Orch-OR theory he developed with the anesthesiologist Stuart Hameroff. Quantum medicine goes out of the scientific game when

it claims that these mechanisms can be controlled from the simple will.

Roger Penrose (1931, English) considers that consciousness results from the reduction of the wave packet and **Elizabeth Rauscher** (1931–2019, American) was a physicist who first became interested in philosophy and then moved on to parapsychology. Elizabeth Rauscher was the cofounder of the Fundamental Fysiks Group in 1975 with George Weissmann to work on quantum foundations. She is co-author with Richard Amoroso of a book on the “holographic anthropic universe” [1072]. They propose a model for creating a scalable quantum computer called “Universal Quantum Computing” based on a theory called “Unified Field Mechanics” which yields no falsifiable prediction and has drawn essentially no engagement from the physics community, placing it in the unsupported claim category rather than in any of the three scientific ones listed earlier. **Richard Amoroso** is Director of the Noetic Advanced Studies Institute in Oakland, California [1073]. Noetics is interested in the links between quantum states and consciousness. This goes well beyond the realm of science with Pragmatic Proof of God [1074].

Steven Weinberg (1933–2021, American), laureate of the Nobel prize in Physics in 1979 for his work on the unification of the weak and electromagnetic forces, thought that philosophy is of little use in quantum physics other than to protect us against the errors of other philosophers [1075]. His view was shared by **Stephen Hawking** (1942–2018, English). See the chapter “Against Philosophy” in “*Dreams of a Final Theory*”, 1992, from Steven Weinberg, which is contradicted by Carlo Rovelli [1076].

Michel Bitbol, trained as a physician and a physicist and now a philosopher of science, is interested in the question of consciousness.

Paavo Pyllkkänen also deals with the place of mind in quantum reality. **Christopher Fields** is using quantum theory and the free energy principle to develop non-reductive explanations. **Amanda Geffer** focuses on subject-object relation and **Jay Garfield** works on the convergence of Mahayana Buddhist metaphysics and quantum mechanics.

The story goes on and on with new interpretations popping up regularly, at least on arXiv (and excluding viXra) [1077–1087]! Some scientists are even trying to explain quantum mechanics with the Bible [1088].

Quantum physics raises other physical-philosophical questions such as does a total vacuum exist? Indeed, quantum physics describes the energy of the vacuum, which would always be crossed by various real and virtual particles. From a practical point of view, it is therefore difficult to create an empty space that is not crossed at all by electromagnetic waves or particles of all kinds. If therefore nothing exists, what was there before the big bang? And let’s not talk about the nature of time, which is still a matter of debate.



3.12.5 Bibliography and notes

These are the bibliographical references and notes for the **Quantum foundations** subsection.

- [117] [The Logic of Quantum Mechanics](#), by Garrett Birkhoff and John Von Neumann, *Annals of Mathematics*, October 1936 (22 pages).
- [123] [Quantum mechanics and reality](#), by Bryce S. DeWitt, 1970 (6 pages)
- [124] [The Many-Worlds Interpretation of Quantum Mechanics](#), by Bryce DeWitt et al., 1973 (146 pages).
- [148] [On the interpretation of measurement in quantum theory](#), by H. Dieter Zeh, *Foundations of Physics*, March 1970 (8 pages)
- [171] [Contexts, Systems and Modalities: a new ontology for quantum mechanics](#), by Alexia Auffèves and Philippe Grangier, *arXiv*, January 2015 (9 pages).
- [391] [The two-spin enigma: from the helium atom to quantum ontology](#), by Philippe Grangier, Alexia Auffèves, Nayla Farouki, Mathias Van Den Bossche, and Olivier Ezratty, *Entropy*, November 2024 (13 pages).
- [488] [The Quantum Measurement Problem: A Review of Recent Trends](#), by Anderson A. Tomaz, Rafael S. Mattos, and Mario Barbatti, *Philosophical Magazine*, December 2025 (38 pages) (preprint on [arXiv](#)).
- [491] [Introduction to Koopman-von Neumann Mechanics](#), by Daniel Piasecki, *arXiv*, December 2021 (123 pages).
- [494] [Heading towards an Algebraic Heisenberg Cut](#), by Mathias Van Den Bossche and Philippe Grangier, *arXiv*, May 2025 (15 pages).
- [825] [Is the universe conscious? It seems impossible until you do the math](#), by Michael Brooks, April 2020.
- [826] [The Mathematical Structure of Integrated Information Theory](#), by Johannes Kleiner and Sean Tull, *arXiv*, February 2020 (22 pages).
- [827] [Bohr and von Neumann on the Universality of Quantum Mechanics: Materials for the History of the Quantum Measurement Process](#), by Federico Laudisa, *arXiv*, May 2024 (29 pages).
- [828] [One Thing Is Certain: Heisenberg's Uncertainty Principle Is Not Dead](#), by Ava Furuta in *Scientific American*, 2012.
- [829] [Einstein, incompleteness, and the epistemic view of quantum states](#), by Nicholas Harrigan and Robert W. Spekkens, *arXiv*, June 2007 (18 pages).
- [830] [In defense of the epistemic view of quantum states: a toy theory](#), by Robert W. Spekkens, *arXiv*, October 2005 (32 pages).
- [831] [On the reality of the quantum state](#), by Matthew F. Pusey, Jonathan Barrett, and Terry Rudolph, *Nature Physics*, November 2012 (3 pages) (preprint on [arXiv](#)).
- [832] [States vs. Changes of States: A Reformulation of the Ontic vs. Epistemic Distinction in Quantum Mechanics](#), by Joanna Luc et al., December 2022 (35 pages).
- [833] [What is ontic and what is epistemic in the Quantum Mechanics of Spin?](#), by Ariel Caticha, *arXiv*, February 2025 (20 pages).
- [834] [A No-Go Theorem for \$\psi\$ -ontic Models? Yes! Response to Criticisms](#), by Gabriele Carcassi, Andrea Oldofredi, and Christine A. Aidala, *arXiv*, December 2024 (8 pages).
- [835] [Is the quantum state real? An extended review of \$\psi\$ -ontology theorems](#), by Matthew Saul Leifer, *Quanta*, November 2014 (89 pages).
- [836] [Why engineers are right to avoid the quantum reality offered by the orthodox theory?](#), by X. Oriols and D. K. Ferry, *arXiv*, February 2024 (6 pages).
- [837] [The Hidden Ontological Variable in Quantum Harmonic Oscillators](#), by Gerard 't Hooft, *arXiv*, October 2024 (12 pages).
- [838] [Quantum probabilities as Bayesian probabilities](#), by Carlton M. Caves, Christopher A. Fuchs, and Rüdiger Schack, *Physical Review A*, January 2002 (6 pages) (preprint on [arXiv](#)).
- [839] [The Future of Quantum Physics](#), by Hans Christian von Baeyer, 2016 (268 pages).
- [840] [Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld](#), by Walther Gerlach and Otto Stern, *Zeitschrift für Physik*, December 1922 (4 pages)
- [841] [Historical Debates over the Physical Reality of the Wave Function](#), by Jacob A. Barandes, *arXiv*, February 2026 (32 pages).
- [842] [A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis](#), by Mordehai Milgrom, *The Astrophysical Journal*, July 1983 (6 pages)
- [843] [On the Physical Untenability of the Standard Notion of Quantum State](#), by Christian de Ronde, *arXiv*, May 2025 (10 pages).
- [844] [Antirealism in sheep's clothing: Wave function realism and scientific realism](#), by Raoni Arroyo and Jonas R. Becker Arenhart, *arXiv*, July 2025 (11 pages).
- [845] [Quantum Physics Is 'Nonsense,' Says Breakthrough Prize Winner Gerard 't Hooft](#), by Lee Billings, *Scientific American*, April 2025.
- [846] [Revisiting the Bohr Model of the Atom through Brownian Motion of the Electron](#), by Vasil Yordanov, *arXiv*, February 2025 (25 pages).
- [847] [The Stochastic-Quantum Correspondence](#), by Jacob A. Barandes, *arXiv*, July 2025 (38 pages).
- [848] [An Einstein-Podolsky-Rosen argument based on weak forms of local realism not falsifiable by GHZ or Bell experiments](#), by Jesse Fulton, *arXiv*, August 2022 (26 pages).
- [849] [Remarks on the mind-body question](#), by Eugene P. Wigner, *The Scientist Speculates*, January 1961 (19 pages)
- [850] [The original Wigner's-Friend scenarios](#), by Jay Lawrence, *arXiv*, September 2025 (6 pages).
- [851] [Wigner's Friend Paradox Revisited](#), by Peter Reichert and Markus Enz, *arXiv*, July 2026 (17 pages).
- [852] [A review and analysis of six extended Wigner's friend arguments](#), by David Schmid, Yilè Yīng, and Matthew Leifer, *arXiv*, September 2024 (32 pages).
- [853] [Mutual Recognition in the Philosophy of Physics: QBism, Phenomenology, Hegel](#), by George Webster, *arXiv*, August 2026 (24 pages).
- [854] [Agreement and Compatibility in Wigner's Friend Paradox](#), by Julio C. F. Silva, B. F. Rizzuti, and Cristhiano Duarte, *arXiv*, June 2026 (26 pages).
- [855] [Intersubjective Agreement about Quantum States Is Unnecessary in QBism](#), by Gino Elia, Jennifer Carter, and Robert Crease, *arXiv*, August 2025 (31 pages).
- [856] [Thought experiments in a quantum computer](#), by Nuriya Nurgalieva, Simon Mathis, Lidia del Rio, and Renato Renner, *arXiv*, September 2022 (10 pages).
- [857] [Wigner's friend scenarios: on what to condition and how to verify the predictions](#), by Flavio Del Santo, Gonzalo Manzano, and Caslav Brukner, *arXiv*, June 2025 (8 pages).

- [858] [Quantum theory cannot consistently describe the use of itself](#), by Daniela Frauchiger and Renato Renner, *Nature Communications*, September 2018.
- [859] [A “thoughtful” Local Friendliness no-go theorem: a prospective experiment with new assumptions to suit](#), by Howard M. Wiseman, Eric G. Cavalcanti, and Eleanor G. Rieffel, *Quantum*, September 2023 (30 pages).
- [860] [A strong no-go theorem on the Wigner’s friend paradox](#), by Kok-Wei Bong, Anibal Utreras-Alarcón, Farzad Ghafari, Yeong-Cherng Liang, Nora Tischler, Eric G. Cavalcanti, Geoff J. Pryde, and Howard M. Wiseman, *Nature Physics*, August 2020 (7 pages).
- [861] [Wigner’s friend’s black hole adventure: an argument for complementarity?](#), by Laurens Wallegghem, *arXiv*, July 2025 (8 pages).
- [862] [Consciousness, Quantum Mechanics, and the Limits of Scientific Objectivism](#), by John B. DeBrot and Christian List, *arXiv*, April 2026 (36 pages).
- [863] [The New Tractatus Program: Quantum Mechanics, Anti-Totalitarian Ontology, and Consciousness](#), by Mikołaj Sienicki and Krzysztof Sienicki, *arXiv*, July 2026 (42 pages).
- [864] [Quantum observers can communicate across multiverse branches](#), by Maria Violaris, *arXiv*, April 2026 (6 pages).
- [865] [Wigner’s Friend as a Circuit: Inter-Branch Communication Witness Benchmarks on Superconducting Quantum Hardware](#), by Christopher Altman, *arXiv*, January 2026 (11 pages).
- [866] [The extended Wigner’s friend, many- and single-worlds and reasoning from observation](#), by Andrew Steane, *arXiv*, May 2026 (31 pages).
- [867] [Is Quantum Mechanics Universal? EWF Experiments and Non Absoluteness of Events](#), by Herve Zwirn, *arXiv*, May 2026 (33 pages).
- [868] [Extended Wigner’s Friend Scenarios with Agent-like Observers on Quantum Computers](#), by Joshua Laux and Eric G. Cavalcanti, *arXiv*, September 2026 (15 pages).
- [869] [Inequivalent ways to apply semi-classical smoothing to a quantum system](#), by Kiarn T. Laverick, Areeya Chantasri, and Howard M. Wiseman, *arXiv*, July 2025 (6 pages).
- [870] [An extended Wigner’s friend no-go theorem inspired by generalized contextuality](#), by Laurens Wallegghem and Lorenzo Catani, *arXiv*, February 2025 (10 pages).
- [871] [Any theory that admits a Wigner’s Friend type multi-agent paradox is logically contextual](#), by Nuriya Nurgalieva and V. Vilasini, *arXiv*, February 2025 (16 pages).
- [872] [The firewall paradox is Wigner’s friend paradox](#), by Ladina Hausmann and Renato Renner, *arXiv*, April 2025 (23 pages).
- [873] [What Do Black Holes Teach Us About Wigner’s Friend?](#), by Emily Adlam, *arXiv*, April 2026 (31 pages).
- [874] [The Interplay between Quantum Contextuality and Wigner Negativity](#), by Pierre-Emmanuel Emeriau, *arXiv*, April 2022 (221 pages).
- [875] [Continuous-variable nonlocality and contextuality](#), by Rui Soares Barbosa, Tom Douce, Pierre-Emmanuel Emeriau, Elham Kashefi, and Shane Mansfield, *arXiv*, April 2022 (44 pages).
- [876] [On the Problem of Hidden Variables in Quantum Mechanics](#), by John S. Bell, *Reviews of Modern Physics*, July 1966 (6 pages) .
- [877] [The Problem of Hidden Variables in Quantum Mechanics](#), by Simon Kochen and Ernst P. Specker, *Journal of Mathematics and Mechanics*, January 1967 (29 pages) .
- [878] [Kochen-Specker Contextuality](#), by Costantino Budroni, Adán Cabello, Otfried Gühne, Matthias Kleinmann, and Jan-Åke Larsson, *arXiv*, January 2023 (68 pages).
- [879] [Coming full circle – A unified framework for Kochen-Specker contextuality](#), by Markus Frembs, *arXiv*, January 2025 (21 pages).
- [880] [Bell-Kochen-Specker theorem: A proof with 18 vectors](#), by Adán Cabello, Jose M. Estebaranz, and Guillermo García Alcaine, *Physics Letters A*, March 1996 (10 pages)  (preprint on *arXiv*).
- [881] [Two fundamental solutions to the rigid Kochen-Specker set problem and the solution to the minimal Kochen-Specker set problem under one assumption](#), by Stefan Trandafir and Adán Cabello, *arXiv*, July 2025 (17 pages).
- [882] [Reply to Pavičić’s ‘Comment on “Optimal conversion of Kochen-Specker sets into bipartite perfect quantum strategies” \(arXiv:2502.13787\)’](#), by Stefan Trandafir and Adán Cabello, *arXiv*, March 2025 (2 pages).
- [883] [Comment on “Optimal conversion of Kochen-Specker sets into bipartite perfect quantum strategies”](#), by Mladen Pavić, *arXiv*, February 2025 (4 pages).
- [884] [Optimal conversion of Kochen-Specker sets into bipartite perfect quantum strategies](#), by Stefan Trandafir and Adán Cabello, *arXiv*, February 2025 (19 pages).
- [885] [The simplest Kochen-Specker set](#), by Adán Cabello, *Physical Review Letters*, November 2025 (7 pages).
- [886] [Statistical signatures of quantum contextuality](#), by Holger F. Hofmann, *arXiv*, June 2024 (19 pages).
- [887] [Measures on the Closed Subspaces of a Hilbert Space](#), by Andrew M. Gleason, *Journal of Mathematics and Mechanics*, January 1957 (9 pages) .
- [888] [Deriving Born’s rule from an Inference to the Best Explanation](#), by Alexia Auffeves and Philippe Grangier, *arXiv*, October 2019 (6 pages).
- [889] [Why qubits are exceptional in Gleason’s theorem: a Bloch-space perspective](#), by Massimiliano Sassoli de Bianchi, *arXiv*, August 2026 (10 pages).
- [890] [Generalised Kochen-Specker Theorem for Finite Non-Deterministic Outcome Assignments](#), by Ravishankar Ramanathan, *arXiv*, October 2024 (14 pages).
- [891] [Contextuality for preparations, transformations, and unsharp measurements](#), by R. W. Spekkens, *arXiv*, March 2005 (16 pages).
- [892] [A structure theorem for generalized-noncontextual ontological models](#), by David Schmid, John H. Selby, Matthew F. Pusey, and Robert W. Spekkens, *arXiv*, March 2024 (41 pages).
- [893] [How contextuality and antidistinguishability are related](#), by Maiyuren Srikumar, Stephen D. Bartlett, and Angela Karanjai, *arXiv*, November 2024 (8 pages).
- [894] [Contextuality Can be Verified with Noncontextual Experiments](#), by Jonathan J. Thio, David R. M. Arvidsson-Shukur, et al., *arXiv*, November 2024 (17 pages).
- [895] [The logical structure of contextuality and nonclassicality](#), by Songyi Liu, Jincheng Wang, et al., *arXiv*, June 2025 (18 pages).
- [896] [Chromatic Quantum Contextuality](#), by Karl Svozil, *arXiv*, April 2025 (7 pages).
- [897] [Context and Complementarity: Formalizing Bohr’s Vision through Many-Valued Contextual Logic](#), by Partha Ghose and Sudip Patra, *arXiv*, May 2025 (10 pages).
- [898] [Complexity of Contextuality](#), by Theodoros Yianni and Farid Shahandeh, *arXiv*, June 2025 (19 pages).
- [899] [Simplicial methods in the resource theory of contextuality](#), by Aziz Kharoof and Cihan Okay, *arXiv*, May 2025 (44 pages).
- [900] [All entangled states can be certified by a violation of noncontextuality inequalities](#), by Yujie Zhang, Robert W. Spekkens, et al., *arXiv*, July 2025 (21 pages).
- [901] [Ruling out nonlinear modifications of quantum theory with contextuality](#), by Ruben Campos Delgado and Martin Plávala, *arXiv*, June 2025 (8 pages).
- [902] [Simulation of quantum contextuality with invasive-measurement maps](#), by Andrea Navoni, Marco G. Genoni, and Andrea Smirne, *arXiv*, July 2025 (5 pages).

- [903] [Empirical Demonstration of Quantum Contextuality on NISQ Computers](#), by Colm Kelleher and Frédéric Holweck, arXiv, May 2025 (8 pages).
- [904] [Contextuality with Pauli observables in cycle scenarios](#), by Raman Choudhary and Rui Soares Barbosa, arXiv, February 2025 (19 pages).
- [905] [Contextuality supplies the 'magic' for quantum computation](#), by Mark Howard, Joel Wallman, Victor Veitch, and Joseph Emerson, Nature, June 2014 (5 pages)  (preprint on [arXiv](#)).
- [906] [Contextuality of Quantum Error-Correcting Codes](#), by Derek Khu, Kishor Bharti, et al., PRX Quantum, January 2026 (29 pages).
- [907] [Quantum retrodiction: foundations and controversies](#), by Stephen M. Barnett, John Jeffers, and David T. Pegg, arXiv, March 2021 (14 pages).
- [908] [The state of a quantum system is not a complete description for retrodiction](#), by Mingxuan Liu, Ge Bai, and Valerio Scarani, arXiv, February 2025 (9 pages).
- [909] [A Retrodictive Approach to Quantum State Smoothing](#), by Mingxuan Liu, Valerio Scarani, Alexia Auffèves, and Kiarn T. Laverick, arXiv, January 2025 (5 pages).
- [910] [Quantum state smoothing as an optimal estimation problem with three different cost functions](#), by Kiarn T. Laverick, Ivonne Guevara, and Howard M. Wiseman, arXiv, June 2021 (7 pages).
- [911] [Retrodiction of measurement outcomes on a single quantum system reveals entanglement with its environment](#), by Julien Pinske and Klaus Mølmer, arXiv, April 2025 (13 pages).
- [912] [Actually, you can't test if quantum uses complex numbers](#), by Craig Gidney, Algassert blog, September 2025.
- [913] [Quantum theory based on real numbers can be experimentally falsified](#), by Marc-Olivier Renou, David Trillo, Mirjam Weilenmann, Thinh P. Le, Armin Tavakoli, Nicolas Gisin, Antonio Acín, and Miguel Navascués, Nature, December 2021 (6 pages).
- [914] [Quantum theory does not need complex numbers](#), by Timothee Hoffreumon and Mischa P. Woods, arXiv, October 2025 (40 pages).
- [915] [Quantum mechanics over real numbers fully reproduces standard quantum theory](#), by Alan C. Maioli, Evaldo M. F. Curado, and Jean-Pierre Gazeau, arXiv, June 2026 (20 pages).
- [916] [Quantum theory based on real numbers cannot be experimentally falsified](#), by Timothée Hoffreumon and Mischa P. Woods, arXiv, March 2026 (60 pages).
- [917] [Comment on "Quantum theory based on real numbers cannot be experimentally falsified": On the compatibility of physical principles with information theory for fermions](#), by Fatemeh Moradi Kalarde, Xiangling Xu, and Marc-Olivier Renou, arXiv, April 2026 (5 pages).
- [918] [A Deflationary Account of Quantum Theory and its Implications for the Complex Numbers](#), by Jacob A. Barandes, arXiv, February 2026 (15 pages).
- [919] [Taxonomy for Physics Beyond Quantum Mechanics](#), by Emily Adlam, Jonte R. Hance, Sabine Hossenfelder, and Tim N. Palmer, arXiv, June 2024 (18 pages).
- [920] [Why does the wavefunction 'collapse' in relational approaches to quantum mechanics?](#), by Emily Adlam, arXiv, April 2026 (27 pages).
- [921] [Another 100 Years of Quantum Interpretation?](#), by Karen Crowther, arXiv, December 2025 (30 pages).
- [922] [Questionable and Unquestionable in Quantum Mechanics](#), by Laszlo E. Szabo, Marton Gomori, and Zalan Gyenis, arXiv, June 2025 (51 pages).
- [923] [Contra multos verbos: On scandals of quantum mechanics](#), by Theodorus Maria Nieuwenhuizen, arXiv, July 2023 (22 pages).
- [924] [The Oxford Handbook of the History of Quantum Interpretations](#), by Olival Freire et al., Oxford University Press, March 2022 (1312 pages) .
- [925] [Physicists disagree wildly on what quantum mechanics says about reality, Nature survey shows](#), by Elizabeth Gibney, Nature, July 2025 .
- [926] [A snapshot of foundational attitudes toward quantum mechanics](#), by Maximilian Schlosshauer, Johannes Kofler, and Anton Zeilinger, Studies in History and Philosophy of Modern Physics, August 2013 (9 pages)  (preprint on [arXiv](#)).
- [927] [On Observation and The Completion of Quantum Mechanics](#), by Matthew F. Brown, arXiv, February 2024 (30 pages).
- [928] [Logical Characterization of Contextual Hidden-Variable Theories based on Quantum Set Theory](#), by Masanao Ozawa, arXiv, November 2023 (12 pages).
- [929] [Who Invented the "Copenhagen Interpretation"? A Study in Mythology](#), by Don Howard, 2004 (15 pages).
- [930] [Copenhagish interpretations of quantum mechanics](#), by David Schmid, Yilè Yīng, and Matthew Leifer, arXiv, July 2025 (16 pages).
- [931] [A Suggested Interpretation of the Quantum Theory in Terms of "Hidden" Variables. I](#), by David Bohm, Physical Review, January 1952 (14 pages) .
- [932] [Harmonizing de Broglie-Bohm's Causal Interpretation with the Copenhagen Interpretation of Quantum Mechanics](#), by A. N. Khondker, arXiv, May 2025 (11 pages).
- [933] [Can Bohmian mechanics be considered complete?](#), by Aurélien Drezet and Arnaud Amblard, arXiv, March 2025 (48 pages).
- [934] [Quantum Physics Without Quantum Philosophy](#), by Detlef Dürr, Sheldon Goldstein, and Nino Zanghì, Springer, 2013 (304 pages) .
- [935] [The trouble with pilot-wave theory: a critical evaluation](#), by Antony Valentini, arXiv, August 2024 (37 pages).
- [936] [Can the ontology of Bohmian mechanics consists only in particles? The PBR theorem says no](#), by Shan Gao, arXiv, January 2025 (25 pages).
- [937] [Information versus Physicality: On the Nature of the Wavefunctions of Quantum Mechanics](#), by C. S. Unnikrishnan, arXiv, June 2025 (17 pages).
- [938] [Pilot-Wave Theories as Hidden Markov Models](#), by Jacob A. Barandes, arXiv, February 2026 (34 pages).
- [939] [De Broglie-Bohm Quantum Mechanics](#), by Antony Valentini, arXiv, September 2024 (31 pages).
- [940] [Bohmian Mechanics: The Physics and Mathematics of Quantum Theory](#), by Detlef Dürr et al., 2009.
- [941] [How does Quantum Uncertainty Emerge from Deterministic Bohmian Mechanics?](#), by Albert Solé, Xavier Oriols, Damiano Marian, and Nino Zanghì, arXiv, October 2016 (16 pages).
- [942] [Revisiting de Broglie's Double-Solution Pilot-Wave Theory with a Lorentz-Covariant Lagrangian Framework](#), by David Darrow and John W. M. Bush, arXiv, August 2024 (31 pages).
- [943] [The Bohmian solution to the problem of time](#), by Ward Struyve, arXiv, July 2023 (13 pages).
- [944] [Why Bohm and Only Bohm?](#), by Jean Bricmont, arXiv, July 2023 (15 pages).
- [945] [Inequivalence of stochastic and Bohmian arrival times in time-of-flight experiments](#), by Pascal Naidon, arXiv, June 2024 (12 pages).
- [946] [Bohmian Mechanics as a Practical Tool](#), by Xabier Oianguren-Asua, Xavier Oriols, et al., arXiv, February 2024 (13 pages).
- [947] [A simple understanding of quantum electrodynamics using Bohmian trajectories: detecting non-ontic photons](#), by Juan José Seoane, Abdelilah Benali, and Xavier Oriols, arXiv, April 2026 (35 pages).

- [948] [Bohmian mechanics: A legitimate hydrodynamic picture for quantum mechanics, and beyond](#), by A. S. Sanz, arXiv, January 2026 (22 pages).
- [949] [Quantum solitodynamics: Non-linear wave mechanics and pilot-wave theory](#), by Aurélien Drezet, arXiv, July 2022 (29 pages).
- [950] [Whence Nonlocality? Removing spooky action at a distance from the de Broglie Bohm pilot-wave theory using a time-symmetric version of de Broglie double solution](#), by Aurélien Drezet, arXiv, December 2023 (50 pages).
- [951] [A time \(anti\)symmetric approach to the double solution theory](#), by Pierre Jamet and Aurélien Drezet, arXiv, October 2024 (33 pages).
- [952] [Arrival-time distributions as a probe of the preferred foliation in relativistic Bohmian mechanics](#), by Arnaud Amblard and Aurélien Drezet, arXiv, April 2026 (39 pages).
- [953] [Geometric Structure of Bell Correlations in Bohmian Mechanics: A Configuration-Space Analysis of EPR Experiments](#), by Tim Dartois, Signe Seidelin, and Aurélien Drezet, arXiv, May 2026 (14 pages).
- [954] [The origin of the Everettian heresy](#), by Stefano Osnaghi et al., 2008 (27 pages)  (preprint on [repositorio.ufba.br](#)).
- [955] [The Multiverse: a Philosophical Introduction](#), by Jeremy Butterfield, arXiv, May 2025 (179 pages).
- [956] [Pure wave mechanics and the notion of worlds](#), by Jeffrey A. Barrett, 2011 (27 pages).
- [957] [Everettian Quantum Mechanics](#), by Stanford Encyclopedia of Philosophy, Stanford Encyclopedia of Philosophy, 1992.
- [958] [Unified Quantum Dynamics: The Emergence of the Born Rule](#), by Martin Weidner, arXiv, May 2025 (71 pages).
- [959] [On playing gods: The fallacy of the many-worlds interpretation](#), by Luis C. Barbado and Flavio Del Santo, arXiv, November 2023 (10 pages).
- [960] [Incoherent? No, Just Decoherent: How Quantum Many Worlds Emerge](#), by Alexander Franklin, arXiv, January 2025 (22 pages).
- [961] [Making Sense of the Many Worlds Interpretation](#), by Stephen Boughn, arXiv, January 2018 (36 pages).
- [962] [The Everything-is-a-Quantum-Wave Interpretation of Quantum Physics](#), by Vlatko Vedral, arXiv, March 2023 (4 pages).
- [963] [Quantum Phenomena Modeled by Interactions between Many Classical Worlds](#), by Michael J. W. Hall et al., PRX, 2014 (17 pages).
- [964] [Everettian Branching in the World and of the World](#), by Nadia Blackshaw, Nick Huggett, and James Ladyman, arXiv, February 2024 (25 pages).
- [965] [Physicists'Views on Scientific Realism](#), by Céline Henne et al., 2024 (41 pages).
- [966] [Defending Many Worlds via Case Discrimination: An Attempt to Showcase the Conceptual Incoherence of Anti-Realist Interpretations and Relational Quantum Mechanics](#), by Michael Huber, April 2023 (25 pages).
- [967] [The Many-Worlds Interpretation of Quantum Mechanics](#), by Lev Vaidman et al., Quantum Reports, March 2023.
- [968] [The Quantum Many-Worlds Interpretation, Simply Told](#), by Brian C. Odom, arXiv, March 2026.
- [969] [Sleeping Beauty in One or Many Worlds: A Defense of the Halfer Position](#), by Jiaxuan Zhang, arXiv, March 2026.
- [970] [Against Many Worlds](#), by Emily Adlam and Jacob A. Barandes, arXiv, July 2026 (22 pages).
- [971] [A one-world interpretation of quantum mechanics](#), by Isaac Layton, Jonathan Oppenheim, and Zachary Weller-Davies, arXiv, October 2025 (19 pages).
- [972] [The Emergent Multiverse: Quantum Theory according to the Everett Interpretation](#), by David Wallace, May 2012 (530 pages) .
- [973] [The Twin-World road to reality in quantum mechanics](#), by Daniel Braun, arXiv, March 2026 (46 pages).
- [974] [The scientific image](#), by Bas Van Fraassen, 1980 (248 pages).
- [975] [Consistent histories and the interpretation of quantum mechanics](#), by Robert B. Griffiths, Journal of Statistical Physics, July 1984 (55 pages) .
- [976] [A Dynamic Histories Interpretation of Quantum Theory](#), by Timothy D. Andersen, arXiv, August 2020 (13 pages).
- [977] [Towards a Deterministic Interpretation of Quantum Mechanics: Insights from Dynamical Systems](#), by Aminur Rahman, arXiv, April 2024 (18 pages).
- [978] [Unified dynamics for microscopic and macroscopic systems](#), by Gian Carlo Ghirardi, Alberto Rimini, and Tullio Weber, Physical Review D, July 1986 (22 pages) .
- [979] [Models of wave-function collapse, underlying theories, and experimental tests](#), by Angelo Bassi, Hendrik Ulbricht, et al., Reviews of Modern Physics, April 2013 (57 pages)  (preprint on [arXiv](#)).
- [980] [Underground test of gravity-related wave function collapse](#), by Sandro Donadi, Angelo Bassi, et al., Nature Physics, January 2021 (5 pages)  (preprint on [arXiv](#)).
- [981] [Rethinking Superdeterminism](#), by Sabine Hossenfelder and Tim Palmer, Front. Phys. May 2020 (13 pages)  (preprint on [arXiv](#)).
- [982] [What does it take to solve the measurement problem?](#), by Jonte R. Hance and Sabine Hossenfelder, arXiv, October 2022 (11 pages).
- [983] [What does the world look like according to superdeterminism](#), by Augustin Baas et al., British Journal for the Philosophy of Science, 2020 (21 pages).
- [984] [A Deterministic Model of Free Will](#), by Tim Palmer, arXiv, April 2025 (21 pages).
- [985] [The burden of Fundamentality: Metaphysical ambiguities and the issue of Superdeterminism](#), by Gabriele Cafiero, Luca Molinari, and Jonte R. Hance, arXiv, August 2026 (24 pages).
- [986] [Relational Quantum Mechanics](#), by Carlo Rovelli, International Journal of Theoretical Physics, September 1996 (42 pages)  (preprint on [arXiv](#)).
- [987] [Relational objectivity in presence of finite quantum resources](#), by Luis C. Barbado and Časlav Brukner, arXiv, February 2025 (14 pages).
- [988] [A Search for Classical Subsystems in Quantum Worlds](#), by Arsalan Adil, Andrew Sornborger, et al., arXiv, May 2024 (18 pages).
- [989] [Comments on 'Comment on Aurélien Drezet's defense of relational quantum mechanics' by Jay Lawrence, Marcin Markiewicz and Marek Żukowski](#), by Aurélien Drezet, arXiv, March 2025 (3 pages).
- [990] [Relative facts do not exist. Relational Quantum Mechanics is Incompatible with Quantum Mechanics. Response to the critique by Aurélien Drezet](#), by Jay Lawrence, Marcin Markiewicz, and Marek Żukowski, arXiv, October 2022 (3 pages).
- [991] [A critical analysis of 'Relative facts do not exist. Relational quantum mechanics is incompatible with quantum mechanics' by Jay Lawrence, Marcin Markiewicz and Marek Żukowski](#), by Aurélien Drezet, arXiv, December 2023 (10 pages).
- [992] [Relative Facts of Relational Quantum Mechanics are Incompatible with Quantum Mechanics](#), by Jay Lawrence, Marcin Markiewicz, and Marek Żukowski, arXiv, May 2023 (12 pages).
- [993] [How to be a Copenhagenistic-QBistic Everettist?](#), by Marcin Wieśniak, arXiv, October 2024 (5 pages).
- [994] [Enaction for QBists](#), by Amanda Geffer, arXiv, December 2024 (23 pages).
- [995] [A no-go result for QBism](#), by Shan Gao, arXiv, January 2025 (8 pages).
- [996] [Redesigning quantum theory](#), by Matthew B. Weiss, arXiv, September 2026 (233 pages).

- [997] [QBism, Polishing Some Points](#), by Christopher A. Fuchs and Blake C. Stacey, arXiv, December 2025 (22 pages).
- [998] [QBism on Locality and Nonlocality](#), by John B. DeBrot, Christopher A. Fuchs, and Rüdiger Schack, arXiv, August 2026 (17 pages).
- [999] [Is the moon there when nobody looks? Reality and the quantum theory](#), by David Mermin, Physics Today, April 1985 (14 pages).
- [1000] [Quantum Wittgenstein - Metaphysical debates in quantum physics don't get at 'truth' – they're nothing but a form of ritual, activity and culture](#), by Timothy Andersen, May 2022.
- [1001] [Zur Quantenmechanik der Stoßvorgänge](#), by Max Born, Zeitschrift für Physik, December 1926 (5 pages) 📄.
- [1002] [The Statistical Interpretation of Quantum Mechanics](#), by Leslie E. Ballentine, Reviews of Modern Physics, October 1970 (24 pages) 📄.
- [1003] [Symmetry of Physical Laws. Part III. Prediction and Retro-diction](#), by Satoshi Watanabe, Reviews of Modern Physics, April 1955 (8 pages) 📄.
- [1004] ["Relative State" Formulation of Quantum Mechanics](#), by Hugh Everett, Reviews of Modern Physics, July 1957 (9 pages) 📄.
- [1005] [Mind, Matter and Quantum Mechanics](#), by Henry P. Stapp, 2009 (303 pages).
- [1006] [The transactional interpretation of quantum mechanics](#), by John G. Cramer, Reviews of Modern Physics, July 1986 (41 pages) 📄.
- [1007] [On gravity's role in quantum state reduction](#), by Roger Penrose, General Relativity and Gravitation, May 1996 (20 pages) 📄.
- [1008] [QBism, the Perimeter of Quantum Bayesianism](#), by Christopher A. Fuchs, arXiv, March 2010 (30 pages).
- [1009] [Violation of Bell's inequalities in a quantum realistic framework](#), by Alexia Auffèves and Philippe Grangier, arXiv, February 2016 (5 pages).
- [1010] [What is quantum in quantum randomness?](#), by Alexia Auffèves and Philippe Grangier, Philosophical Transactions of the Royal Society A, April 2018 (9 pages) 📄 (preprint on arXiv).
- [1011] [Extracontextuality and Extravalence in Quantum Mechanics](#), by Alexia Auffèves and Philippe Grangier, Philosophical Transactions of the Royal Society A, May 2018 (7 pages) 📄 (preprint on arXiv).
- [1012] [A generic model for quantum measurements](#), by Alexia Auffèves and Philippe Grangier, arXiv, July 2019 (8 pages).
- [1013] [Contextual inferences, nonlocality, and the incompleteness of quantum mechanics](#), by Philippe Grangier, arXiv, January 2022 (7 pages).
- [1014] [Revisiting Quantum Contextuality in an Algebraic Framework](#), by Mathias Van Den Bossche and Philippe Grangier, arXiv, April 2023 (8 pages).
- [1015] [Kolmogorovian Censorship, Predictive Incompleteness, and the locality loophole in Bell experiments](#), by Philippe Grangier, arXiv, March 2025 (4 pages).
- [1016] [The Quantum Mechanics Conundrum](#), by Gennaro Auletta, 2019 (879 pages).
- [1017] [A categorical semantics of quantum protocols](#), by Samson Abramsky and Bob Coecke, Proceedings of the 19th IEEE Symposium on Logic in Computer Science, July 2004 (21 pages).
- [1018] [Are Words the Quanta of Human Language? Extending the Domain of Quantum Cognition](#), by Diederik Aerts and Lester Beltran, arXiv, December 2021 (27 pages).
- [1019] [From Quantum Cognition to Conceptuality Interpretation I: Tracing the Brussels Group's Intellectual Journey](#), by Diederik Aerts, Massimiliano Sassoli de Bianchi, and Sandro Sozzo, arXiv, November 2024 (15 pages).
- [1020] [The cognitive triple-slit experiment](#), by Luca Sassoli de Bianchi and Massimiliano Sassoli de Bianchi, arXiv, June 2025 (23 pages).
- [1021] [From Quantum Cognition to Conceptuality Interpretation II: Unraveling the Quantum Mysteries](#), by Diederik Aerts, Massimiliano Sassoli de Bianchi, and Sandro Sozzo, arXiv, January 2025 (19 pages).
- [1022] ["What is a Thing?": Topos Theory in the Foundations of Physics](#), by Andreas Doering and Chris Isham, arXiv, March 2008 (212 pages).
- [1023] [The physical interpretation of daseinisation](#), by Andreas Doering, arXiv, April 2010 (39 pages).
- [1024] [Review of the Topos Approach to Quantum Theory](#), by Cecilia Flori, arXiv, June 2011 (24 pages).
- [1025] [Quantum Darwinism, Classical Reality, and the Randomness of Quantum Jumps](#), by Wojciech H. Zurek, arXiv, December 2014 (8 pages).
- [1026] [High-dimensional monitoring and the emergence of realism via multiple observers](#), by Alexandre C. Orthey Jr., Pedro R. Dieguez, Owidiusz Makuta, and Remigiusz Augusiak, arXiv, February 2025 (10 pages).
- [1027] [Classifying two-body Hamiltonians for Quantum Darwinism](#), by Emery Doucet and Sebastian Deffner, arXiv, January 2025 (24 pages).
- [1028] [Quantum Darwinism-encoding transitions on expanding trees](#), by Benoît Ferté and Xiangyu Cao, arXiv, March 2024 (28 pages).
- [1029] [Consensus About Classical Reality in a Quantum Universe](#), by Akram Touil, Bin Yan, and Wojciech H. Zurek, arXiv, March 2025 (4 pages).
- [1030] [Demystifying Objectivity with Operator Algebra Quantum Error Correction](#), by Marin Girard, Gong Cheng, and ChunJun Cao, arXiv, June 2026 (5 pages).
- [1031] [What Kind of World Supports Darwinian Evolution? Quantum Foundational Options](#), by Partha Ghose, arXiv, February 2026 (8 pages).
- [1032] [From compatibility of measurements to exploring Quantum Darwinism on NISQ](#), by Emery Doucet and Sebastian Deffner, arXiv, January 2026 (9 pages).
- [1033] [Decoherence, einselection, and the quantum origins of the classical](#), by Wojciech Hubert Zurek, Reviews of Modern Physics, May 2003 (61 pages) 📄 (preprint on arXiv).
- [1034] [Natural Probability](#), by Brett Parker, arXiv, December 2024 (50 pages).
- [1035] [Observation of quantum Darwinism and the origin of classicality with superconducting circuits](#), by Zitian Zhu, Rubem Mondaini, et al., Science Advances, August 2025 (9 pages).
- [1036] [A paradox in the quantum-mechanical treatment of destructive measurements on photons](#), by Marcoen J.T.F. Cabbolet and Yves Caudano, arXiv, May 2025 (6 pages).
- [1037] [Hamiltonian Systems and Transformation in Hilbert Space](#), by Bernard O. Koopman, Proceedings of the National Academy of Sciences, May 1931 (4 pages) 📄 (preprint on ncbi.nlm.nih.gov).
- [1038] [Zur Operatorenmethode in der klassischen Mechanik](#), by John von Neumann, Annals of Mathematics, July 1932 (56 pages) 📄.
- [1039] [An Algebraic Approach to Koopman Classical Mechanics](#), by Peter Morgan, arXiv, January 2020 (123 pages).
- [1040] [Implementing the Koopman-von Neumann approach on continuous-variable photonic quantum computers](#), by Xinfeng Gao, Olivier Pfister, and Stefan Bekiranov, arXiv, December 2025 (9 pages).
- [1041] [Solutions of Koopman-von Neumann equations, their superpositions, orthogonality and uncertainties](#), by Mustafa Amin and Mark A. Walton, arXiv, December 2025 (17 pages).

- [1042] [Bosonic Encodings for Hermite-Galerkin Discretizations of High-Dimensional PDEs and Bayesian Inverse Problems](#), by Alice Barthe, arXiv, August 2026 (33 pages).
- [1043] [Classical Systems in Quantum Mechanics](#), by Pavel Bóna, arXiv, July 2020 (230 pages).
- [1044] [Interpreting Quantum Entanglement: Steps Towards Coherentist Quantum Mechanics](#), by Claudio Calosi et al., 2018 (51 pages).
- [1045] [Epistemic Horizons From Deterministic Laws: Lessons From a Nomic Toy Theory](#), by Johannes Fankhauser, Tomáš Gonda, and Gemma De les Coves, arXiv, June 2024 (11 pages).
- [1046] [An Indeterminacy-based Ontology for Quantum Theory](#), by Francisco Pipa, arXiv, August 2024 (52 pages).
- [1047] [Relational Quantum Dynamics \(RQD\): An Informational Ontology](#), by Arash E. Zaghi, arXiv, May 2025 (29 pages).
- [1048] [Constructor Theory](#), by David Deutsch, arXiv, January 2013 (40 pages).
- [1049] [The Science of Can and Can't: A Physicist's Journey through the Land of Counterfactuals](#), by Chiara Marletto, May 2021.
- [1050] [Counterfactuals in Macroscopic Quantum Physics: Irreversibility, Measurement and Locality](#), by Maria Violaris, arXiv, May 2025 (179 pages).
- [1051] [Constructor Theory as Process Theory](#), by Stefano Gogioso, Bob Coecke, et al., arXiv, December 2023 (15 pages).
- [1052] [Collapse-based models for gravity do not violate the Entanglement-Based Witness of non-classicality](#), by Tianfeng Feng, Vlatko Vedral, and Chiara Marletto, arXiv, March 2025 (5 pages).
- [1053] [Classical gravity cannot mediate entanglement](#), by Chiara Marletto, Jonathan Oppenheim, Vlatko Vedral, and Elizabeth Wilson, arXiv, November 2025 (6 pages).
- [1054] [Classical Gravity Cannot Mediate Entanglement by Local Means](#), by Chiara Marletto and Vlatko Vedral, arXiv, October 2025 (2 pages).
- [1055] [Interference in complex canonical variables is not quantum](#), by Chiara Marletto and Vlatko Vedral, arXiv, May 2025 (3 pages).
- [1056] [Constructor theory of time](#), by David Deutsch and Chiara Marletto, arXiv, May 2025 (13 pages).
- [1057] [Bose-Marletto-Vedral experiment without observable space-time superpositions](#), by Nicetu Tibau Vidal, Chiara Marletto, Vlatko Vedral, and Giulio Chiribella, arXiv, June 2025 (2 pages).
- [1058] [Tests of constructor theory](#), by Chiara Marletto, David Deutsch, and Vlatko Vedral, arXiv, June 2026 (35 pages).
- [1059] [Universal quantum theory from dynamical consistency](#), by Chiara Marletto and Vlatko Vedral, arXiv, August 2026 (5 pages).
- [1060] [Quantum Measurement Without Collapse or Many Worlds: The Branched Hilbert Subspace Interpretation](#), by Xing M. Wang, arXiv, June 2025 (18 pages).
- [1061] [Einstein's Electron and Local Branching: Unitarity Does not Require Many-Worlds](#), by Xing M. Wang, arXiv, July 2025 (11 pages).
- [1062] [Algorithmic idealism: what should you believe to experience next?](#), by Markus P. Mueller, arXiv, July 2025 (32 pages).
- [1063] [Algorithmic Idealism III: "Algorithmic State" Formulation of Quantum Mechanics](#), by Krzysztof Sienicki, arXiv, January 2025 (29 pages).
- [1064] [Algorithmic Idealism II: Reassessment of Competing Theories](#), by Krzysztof Sienicki, arXiv, December 2024 (20 pages).
- [1065] [Algorithmic Idealism I: Reconceptualizing Reality Through Information and Experience](#), by Krzysztof Sienicki, arXiv, December 2024 (17 pages).
- [1066] [Quantum Reconstructions as Stepping Stones toward Psi-Doxastic Interpretations?](#), by Philipp Berghofer, arXiv, October 2024 (36 pages).
- [1067] [Defending the Quantum Reconstruction Program](#), by Philipp Berghofer, arXiv, October 2024 (37 pages).
- [1068] [Quantum Probabilities Are Objective Degrees of Epistemic Justification](#), by Philipp Berghofer, arXiv, October 2024 (23 pages).
- [1069] [Super Quantum Mechanics](#), by Mikhail Gennadievich Belov, Vladislav Gennadievich Malyshev, et al., arXiv, January 2025 (13 pages).
- [1070] [Time Symmetry, Retrocausality, and Emergent Collapse: The Tlalpan Interpretation of Quantum Mechanics](#), by Alejandro Frank, arXiv, August 2025 (9 pages).
- [1071] [Perspectival Objectivity or: How I Learned to Stop Worrying and Love Observer-Dependent Reality](#), by Peter W. Evans, 2020 (16 pages).
- [1072] [The Holographic Anthropic Universe](#), by Richard L. Amoroso et al., 2009 (510 pages).
- [1073] [Brief Primer on the Fundaments of Quantum Computing](#), by Richard L. Amoroso, 2017 (140 pages).
- [1074] [Pragmatic Proof of God](#), by Richard L. Amoroso et al., August 2017.
- [1075] [The Trouble with Quantum Mechanics](#), by Steve Weinberg, 2016.
- [1076] [Physics Needs Philosophy - Philosophy Needs Physics](#), by Carlo Rovelli, 2023.
- [1077] [An alternative foundation of quantum theory](#), by Inge S. Helland, arXiv, November 2023 (49 pages).
- [1078] [A new approach towards quantum foundation and some consequences](#), by Inge S. Helland, arXiv, February 2025 (20 pages).
- [1079] [The final version of a recent approach towards quantum foundation](#), by Inge S. Helland, arXiv, April 2026 (16 pages).
- [1080] [A Theory of Inaccessible Information](#), by Jacopo Surace, arXiv, September 2024 (17 pages).
- [1081] [Is the Statistical Interpretation of Quantum Mechanics \$\psi\$ -Ontic or \$\psi\$ -Epistemic?](#), by Mario Hubert, arXiv, November 2022 (26 pages).
- [1082] [The solution to the "Einstein-Podolsky-Rosen paradox"](#), by Roman Schnabel, arXiv, February 2023 (8 pages).
- [1083] [Epistemic-Pragmatist Interpretations of Quantum Mechanics: A Comparative Assessment](#), by Ali Barzegar and Daniele Oriti, arXiv, February 2025 (40 pages).
- [1084] [Understanding Quantum Theory: An Operational Reconstructive Approach](#), by Philip Goyal, arXiv, April 2026 (30 pages).
- [1085] [The Transformation-Response Framework: An Operational Reformulation of Quantum Mechanics](#), by Meng-Jun Hu, arXiv, June 2026 (31 pages).
- [1086] [Quantum Processes under Epistemic Constraints](#), by Varun Immanuel, arXiv, August 2026 (188 pages).
- [1087] [Systematizing the Interpretation of Quantum Theory via Reconstruction](#), by Philip Goyal, arXiv, December 2025 (31 pages).
- [1088] [A biblical view restores reality to quantum mechanics](#), by D. Russell Humphreys, Journal of Creation, 2020 (6 pages).

This **Quantum foundations** subsection contains 273 bibliographical references and notes containing at least 11,213 pages.

3.13 Quantum physics 101 key takeaways

- Quantum physics is based on a set of postulates and a linear algebra mathematical formalism. Surprisingly, there are many variations of these postulates. There is not a single bible or reference for these, illustrating the diversity of pedagogies and interpretations in quantum physics. One big underlying question is “what is reality?”. But although deemed incomplete, the theory has been validated by an incredible number of experiments and to extreme precision.
- Quantum physics describes the behavior of matter and light at the nanoscale level, but it can in some conditions extend to larger objects like molecules or even artificial atoms like superconducting currents, Bose Einstein condensates and quantum dots. It deals not only with atoms, electrons and photons which are used in quantum information technologies but also with all elementary particles from the standard model (quarks, ...). However, we don't use or need to look at this level of granularity in most quantum technologies.
- Quantumness comes from the quantization of many properties of light and matter that can take only discrete values, from the wave-particle duality of massive particles such as atoms and electrons and massless particles such as photons, and from phenomena such as superposition and entanglement. Atoms, electrons, nucleons and photons can be described by quantum numbers associated with particular observables, while quantum systems can also have continuous observables such as position and momentum. Schrödinger's cat is a thought experiment about measurement, entanglement and macroscopic definiteness. Macroscopic quantum superpositions are not forbidden merely by an object's size, but environmental decoherence suppresses their observable coherence extremely rapidly in ordinary conditions.
- The Heisenberg uncertainty principle states that the statistical uncertainties of incompatible observables obey uncertainty relations, for example $\Delta x \Delta p \geq \hbar/2$ for position and momentum. Quantum squeezing can reduce the uncertainty of one observable at the cost of increased uncertainty in an incompatible observable. Optical squeezing is used in gravitational-wave interferometers such as LIGO and in other areas of quantum photonics and sensing. The uncertainty principle is not restricted to continuous observables: incompatible discrete observables also obey uncertainty relations, whereas commuting observables can in principle be measured jointly. This part also looks at the role of measurement and symmetries in quantum physics.
- Quantum matter and fluids are showing up with composite elements associating light and matter, or with superfluidity and superconductivity where boson quantum objects can behave like a single quantum object. You find there a wealth of strange phenomena such as skyrmions, magnons, topological insulators and quantum batteries. They could lead to a new chapter in the second quantum revolution.
- Quantum physics also explains weird effects like vacuum quantum fluctuation, although it doesn't violate the second principle of thermodynamics, nor can it lead to the creation of some free energy sources.
- Most quantum physics phenomena described in this section have or will have some use cases in quantum information science and technologies. Sometimes, the same quantum objects and properties like cold atoms can be used across quantum computing, quantum sensing and quantum communications.
- Quantum foundations is the branch philosophy of science that aims to build some understanding of the real world. Quantum physics' formalism is difficult to associate with the principles of reality usually applicable in classical physics. While classical physics understanding has historically been associated with an ontology with objects' position and motion enabling the prediction of phenomena such as the motion of planets. Quantum physics lacks such an ontology describing the physical world. Beyond the canonical Copenhagen interpretation (psi and the wave equation), many scientists tried to create such ontologies, and the debate is still raging among them.

4 Gate-based quantum computing

As a computer scientist, you may have skipped all the previous parts to get here right away. One can indeed understand how quantum computers operate without delving too deeply into quantum physics beyond grasping its basic mechanisms. I still believe that quantum physics and quantum computing hardware are domains of interest to understand what is possible now and will be possible in the future as a quantum software developer. Some mathematical knowledge is, however, required on trigonometry and linear algebra, including vectors, matrices, and complex numbers which are covered here. Here, we'll learn the basics of qubits, qubit registers, quantum gates, understanding what is a quantum circuit and the lifecycle of a qubit from initialization to measurement. Later in this book, after the quantum computing hardware related parts, we'll look at quantum algorithms, software tools and applications. This part is a bit like the “macro-assembly” equivalent for quantum computing while the other parts are about algorithms, programming, software patterns, development tools and applications.

The first basic element of a quantum computer is its inevitable qubit. You've probably already heard about this object having “simultaneously” the values 0 and 1. It has not. A qubit can be prepared in a coherent superposition $\alpha|0\rangle + \beta|1\rangle$, whose measurement statistics depend on the amplitudes and on their relative phase. You may also know that a set of N qubits creates an exponential 2^N dimensional state space that would explain the power of quantum computing. An N qubits pure state does have 2^N computational-basis amplitudes, but a Hilbert-space dimension explains no useful computing power alone. A speedup requires an algorithm that controls interferences and extracts a value of interest efficiently. Unfortunately, most explanations usually stop there, and you then end up wondering how this can run real useful computation. What data comes in and out of a quantum computer? How is it programmed? How do you feed it with data and code? Where is it useful? Why is it potentially faster than classical computing? What is the nature of the “quantum parallelism” exhibited by quantum computers? This book is there to provide you with some educated answers to all these critical questions.

Adopting a bottom-up approach, we describe in this part the logic and mathematics behind qubit manipulation, qubit registers, quantum circuits, quantum gates and measurement. Surprisingly, the term qubit, for ‘quantum’ and ‘bit’, appeared in 1995, not long after, and not before, the creation of some foundational quantum algorithms like Deutsch-Jozsa and Shor integer factoring [1089].

When possible, we draw parallels with classical computing. In this step by step journey, we later look at quantum computer engineering and hardware, describe the complete architecture of a superconducting qubit quantum computer as an example, look at the generic aspects of error mitigation and correction, at the energetics of quantum computing and quantum memories. Other parts later in the book cover the fields of quantum algorithms, software tools and use cases. You can also

choose to start there, in a top-to-bottom approach, and come back here for the details. That's your choice.

4.1 Linear algebra

Quantum physics and computing require some understanding of a whole bunch of mathematical concepts which are mainly related to linear algebra. They are associated with a mathematical formalism describing quantum phenomena. It is also the cornerstone of quantum physics postulates, already covered in an earlier section. It is also essential to understand how qubit, quantum gates, quantum algorithms and quantum measurement operate.

4.1.1 Linearity

Linear algebra is the branch of mathematics using vector spaces, matrices and linear transformations. Quantum physics and computing also rely extensively on complex numbers. Gerolamo Cardano (1501-1576, Italian) explicitly encountered square roots of negative numbers in his 1545 *Ars Magna*, and Rafael Bombelli (1526-1572, Italian) later developed systematic rules for manipulating them in *L'Algebra*, published in 1572 [1090]. Complex numbers are used throughout quantum physics and computing.

A phenomenon is linear if its effects are proportional to its causes. This translates into the verification of two simple equations pertaining to homogeneity and additivity.

In a one-dimensional space, a classic example of a linear function is $f: x \mapsto ax$. A polynomial function of the type $f: x \mapsto ax^2 + b$ is obviously not linear because it evolves non-proportionally to x . Even $f: x \mapsto ax + b$ is not linear, and for the same reason.

As already defined, an observable is represented by a Hermitian operator whose eigenvalues are the possible outcomes of an ideal projective measurement. For a

qubit, a computational-basis measurement produces a classical 0 or 1 and projects the state onto $|0\rangle$ or $|1\rangle$. An immediate repetition of the same ideal projective measurement gives the same result only if there is no intervening evolution or disturbance. Other measurement bases can be used by choosing a different observable or by rotating the state before measurement. Photon polarization measurements are a common example.

Q Defining linearity

Linearity is based on two mathematical properties applied to a map $f : V \rightarrow W$ between vector spaces over the same field, which is usually $\mathbb{C}$ in quantum mechanics:

Homogeneity

$f(\lambda x) = \lambda f(x)$ for all $x \in V$ and scalars λ

Additivity

$f(x + y) = f(x) + f(y)$ for all $x, y \in V$.

Together these properties are equivalently written as $f(\alpha x + \beta y) = \alpha f(x) + \beta f(y)$.

4.1.2 Hilbert spaces and orthonormal bases

A quantum state of a single or several quantum objects can be described by a vector in a Hilbert space. A qubit state is represented in a two-dimensional orthonormal space formed with the basis states vectors $|0\rangle$ and $|1\rangle$. A pure state is a normalized ray in a complex Hilbert space, whose inner product defines norms, overlaps and orthogonality, a global phase leaving the same ray unchanged. A complex number is defined as $a+ib$ where a and b are real and $i^2 = -1$ (Figure 170).

Complex numbers are central to quantum mechanics because quantum states are represented by complex probability amplitudes. Their moduli determine measurement probabilities through the Born rule, while their relative phases control interference and unitary dynamics.

With qubits, the state is represented by the complex amplitudes α and β associated with $|0\rangle$ and $|1\rangle$, with the normalization $|\alpha|^2 + |\beta|^2 = 1$. This linear combination describes a pure single-qubit superposition.

A qubit is an effective two-dimensional Hilbert-space degree of freedom selected from the generally larger Hilbert space of its physical carrier. For example, a position wave function may live in an infinite-dimensional space, while two selected atomic levels, spin states or polarization states can encode a qubit. Photonic qubits are quantum states of the electromagnetic field. Classical Maxwell equations describe classical field propagation, while single-photon states and their measure-

ments require quantum electrodynamics or an appropriate quantum-optical model.




Figure 170: Complex number explained by geometry and trigonometry. r = amplitude, modulus, norm, θ = phase angle, Euler's formula: $e^{i\theta} = \cos \theta + i \sin \theta$, phase angles add up like: $e^{i(\theta_1 + \theta_2)} = e^{i\theta_1} \cdot e^{i\theta_2}$.

An orthonormal basis of a vector space consists of base vectors which are all mathematically orthogonal with each other and whose length is 1. In the representation of a qubit state, the most common orthonormal basis is made of the states $|0\rangle$ and $|1\rangle$.

Other orthonormal reference basis can be used for measurement, particularly with photons, and polarization references different from the starting reference ($0^\circ/90^\circ$ then $45^\circ/135^\circ$, obtained with rotating a simple polarizer).

Another example of an orthonormal basis is the pair of states located on the Bloch-sphere x-axis and represented by $|+\rangle$ and $|-\rangle$. These are the eigenstates of the Pauli X operator, usually called the X-basis states:

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \quad |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

4.1.3 Dirac Notation

In Dirac notation, a quantum object state is represented by $|\psi\rangle$, the **ket** of quantum state Ψ . The **bra** of the same state vector, represented by $\langle\Psi|$, is the conjugate transpose, also named Hermitian adjoint, of the “ket”.

$$\text{bra: } \langle\Psi| = [\bar{\alpha}, \bar{\beta}] \quad \text{ket: } |\Psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}$$

It is the “horizontal” vector $[\bar{\alpha}, \bar{\beta}]$ where $\bar{\alpha}$ and $\bar{\beta}$ are the conjugates of α and β , inverting the sign of the imaginary part of the number ($-i$ instead of $+i$, or the opposite). Here is a complex number conjugate example:

$$(1 + i)^* = 1 - i$$

The **scalar product** of two qubits $\langle \Psi_1 | \Psi_2 \rangle$ is the mathematical projection of the state vector $|\Psi_2\rangle$ onto the vector $|\Psi_1\rangle$. This yields a complex number. When the vectors are orthogonal, the scalar product is equal to 0. When the two vectors are identical, $\langle \Psi | \Psi \rangle$ is Ψ 's squared norm, which equals 1 for a normalized state vector. A scalar product is also named an inner product:

$$\langle A | B \rangle = \sum_{i=1}^n \bar{a}_i b_i$$

An **inner product** is a generalization of a dot vector product applied to complex number vectors, which computes vector similarity, according to examples below:

$$\begin{aligned} \langle \Psi_1 | \Psi_2 \rangle &= [\bar{\alpha}_1, \bar{\beta}_1] \times \begin{bmatrix} \alpha_2 \\ \beta_2 \end{bmatrix} \\ &= \bar{\alpha}_1 \alpha_2 + \bar{\beta}_1 \beta_2 \end{aligned}$$

And on using the same Ψ :

$$\begin{aligned} \langle \Psi | \Psi \rangle &= [\bar{\alpha}, \bar{\beta}] \times \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \\ &= |\alpha|^2 + |\beta|^2 \end{aligned}$$

The **outer product** of two vectors representing a qubit, one in bra and the other in ket, gives an operator or density matrix which is a 2×2 matrix:

When the bra corresponds to the transconjugate of the ket, it is a density operator of a pure state. This notion of density operator will then be extended to a combination of qubits.

$$|\Psi\rangle \langle \Psi| = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \times [\bar{\alpha}, \bar{\beta}] = \begin{bmatrix} \alpha \bar{\alpha} & \alpha \bar{\beta} \\ \beta \bar{\alpha} & \beta \bar{\beta} \end{bmatrix}$$

What are the use cases of this Dirac notation? It is particularly helpful for manipulating quantum states, to simplify tensor products representations and with measurement.

4.1.4 Eigenstuff

We also need to define the notions of **eigenvector**, **eigenvalue**, **eigenstate** and **eigenspace** which are often used in quantum mechanics and quantum computing as well as in machine learning, particularly in dimension reduction algorithms such as PCA (Principal

Components Analysis). These notions allow to define the structure of certain square matrices [1091].

For a square matrix A , an eigenvector x or eigenvalue of A is a vector that verifies the equation $Ax = \lambda x$, λ being a complex number called eigenvalue.

These eigenvectors have the particularity of changing only by a scalar factor when multiplied by the matrix A . For an eigenvalue λ , the associated eigenspace is the set of vectors x satisfying $Ax = \lambda x$. The eigenvalues are roots of the characteristic equation $\det(A - \lambda I) = 0$, where I is the identity matrix. For an $n \times n$ matrix, the characteristic polynomial has degree n [1092].

The reference eigenvectors of a matrix A allow to reconstitute an orthonormal space linked to the matrix. For example, a projection matrix in a 3D plane will have as main eigenvectors two orthogonal vectors located in the plane and one vector orthogonal to the plane. This multiplication gives λx with λ being non-zero if the eigenvector is in the plane in question and 0 if the vector is orthogonal to the plane [1093]. A matrix A can be that of a quantum gate. An eigenvector $|\psi\rangle$ of a unitary gate U satisfies $U|\psi\rangle = \lambda|\psi\rangle$. Its ray is unchanged, while the ket can acquire the eigenvalue phase λ .

This is easy to imagine for the S gate, phase change, which we will see later. The $|0\rangle$ and $|1\rangle$ kets being in the rotation axis, they are not modified by it.

They are thus eigenvectors of the $S = \text{diag}(1, i)$ gate and the corresponding eigenvalues are 1 and i . More generally, every eigenvalue of a unitary matrix has modulus 1, so it can be written $e^{i\theta}$. Eigenvalues are not restricted to +1 and -1.

Finding the eigenvectors and eigenvalues of a matrix A supports its diagonalization when A has a complete set of linearly independent eigenvectors. Eigenvalues exist for every complex square matrix, but a defective matrix has no eigenbasis. Hermitian and unitary matrices commonly used in quantum physics are all non-defective and diagonalizable. The diagonalization of a square matrix consists of finding the matrix which will multiply it to transform it into a matrix filled only in its diagonal. A matrix A is diagonalizable if we can find a matrix P and a diagonal matrix D such that $P^{-1}AP = D$ (P^{-1} being the inverse matrix of P , such that $P^{-1}P = PP^{-1} = I$, I being the matrix identity with 1's in the diagonal and 0's elsewhere). A square matrix of dimension n is diagonalizable if it has n mutually independent eigenvectors. The diagonalized matrix diagonal contains the eigenvalues λ_i of the origin matrix, with $i=1$ to N being the size of the matrix.

A normal operator, and in particular a Hermitian operator or density matrix, admits a spectral decomposition of the form $A = \sum_i \lambda_i |i\rangle \langle i|$, where the $|i\rangle$ form an orthonormal eigenbasis. This is an operator decomposition, not a decomposition of a pure state vector.




Figure 171: An idealized view of a Gaussian wave packet. For a single photon, the packet is a coherent superposition of frequency modes, not an addition of several photons at different wavelengths.

A wave packet can instead be written as a coherent superposition of frequency or momentum modes. For a single photon, a Gaussian spectral envelope can define one photonic wave-packet mode containing one quantum of excitation. For a classical femtosecond pulse, the corresponding field envelope is similarly built from a coherent bandwidth of optical frequencies. A pulse whose instantaneous carrier frequency varies in time is called chirped.

And the eigenstates? This is another name given to eigenvectors, but by physicists!

4.1.5 Tensor products

The tensor product of two vectors of dimension m and n gives a vector of dimension $m \times n$ while the tensor product of a matrix of dimension $m \times n$ by a matrix of dimension $k \times l$ will give a matrix of dimension $mk \times nl$. Tensor products use the sign $\otimes$ (Figure 173).

Tensor products are used to compute “manually” the state of quantum registers containing several unentangled qubits. The state of a register of N non-entangled qubits is the tensor product of these N qubits represented by their vertical ket vector.

This gives a ket with 2^N complex amplitudes, one for each computational-basis bit string. A general register state is $|\Psi\rangle = \sum_x \alpha_x |x\rangle$, with the Born normalization $\sum_x |\alpha_x|^2 = 1$. For a separable register these amplitudes factorize as a tensor product of the individual qubit state vectors. Entangled states generally cannot be factorized this way.

4.1.6 Entanglement

Quantum states are separable when they are mathematically the result of the tensor product of each of the pure states that compose it. But these values can be assembled linearly to create another quantum state, modulo a normalization rule. This combines several vectors resulting from tensor products. These combinations can become inseparable. That’s when entanglement comes into play. An entangled state of two or more qubits occurs when it cannot be factorized as the tensor product of two pure states. In other words, it cannot be the combination of independent qubits. The qubits become dependent.

This is demonstrated mathematically for the states $|00\rangle$ and $|11\rangle$ of a register of two qubits.

In these pairs, the measurement of the value of one of the qubits determines that of the other, here identical. The creation of such entangled pairs of qubits requires preparation operations like using a combination of Hadamard and CNOT gates.

Two qubits placed side by side are not magically entangled! The pair used in the example can be generated by two quantum gates, an H gate (Hadamard) and a CNOT gate (Figure 172). The Bell pair combines the states $|00\rangle$ and $|11\rangle$.




Figure 172: A Bell pair constructed with a Hadamard and a CNOT gate.

We will define this CNOT gate later. Each local Z outcome is random while the joint outcomes are perfectly correlated. The two qubits do not carry preexisting values. The entangled state also predicts basis-dependent correlations that no pair of independent local pure states can reproduce.

An entangled EPR pair can’t be a tensor product of two qubits $|\Psi_1\rangle$ and $|\Psi_2\rangle$

An EPR pair is an entangled state. For example, consider the Bell state:

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle)$$

Suppose this were a product state of two individual qubits:

$$|\Phi^+\rangle = |\Psi_1\rangle \otimes |\Psi_2\rangle$$

Let us assume general single-qubit pure states:

$$|\Psi_1\rangle = \alpha |0\rangle + \beta |1\rangle, \quad |\Psi_2\rangle = \gamma |0\rangle + \delta |1\rangle$$

Then the tensor product becomes:

$$|\Psi_1\rangle \otimes |\Psi_2\rangle = \alpha\gamma |00\rangle + \alpha\delta |01\rangle + \beta\gamma |10\rangle + \beta\delta |11\rangle$$

For this to match $|\Phi^+\rangle$, we would need:

$$\alpha\gamma = \frac{1}{\sqrt{2}}, \quad \beta\delta = \frac{1}{\sqrt{2}}, \quad \alpha\delta = \beta\gamma = 0$$

This leads to a contradiction since all four coefficients cannot satisfy both the amplitude values and the vanishing cross terms. Hence, $|\Phi^+\rangle$ cannot be written as a product of two qubit states.




Figure 173: Tensor products construction. © Olivier Ezratty, 2020-2025.

In a gate-based circuit, entangling multi-qubit gates are the standard way to create entanglement from initially separable qubits. Not every multi-qubit gate is entangling, and an entangling gate does not entangle every possible input state. Entanglement can also be generated by physical interactions, joint measurements with conditioning, or engineered dissipation.

A so-called **GHZ** state (for Greenberger-Horne-Zeilinger, distinguishable from GHz frequencies with a capital Z) with three (or more) entangled qubits is superposing the states $|000\rangle$ and $|111\rangle$. It is a generalization of the 2-qubit Bell state $(|00\rangle + |11\rangle)/\sqrt{2}$. A GHZ is usually prepared with a Hadamard gate and two consecutive CNOTs (Figure 174).




Figure 174: A three-qubit GHZ state constructed with an Hadamard and two CNOT gates.

These pairs of Bell and GHZ states are used in error correction codes as well as in telecommunications, among other things.

Another typical genuinely multipartite entangled state is the **W state**. Together with GHZ states, it represents one of the two inequivalent classes of three-qubit entanglement under stochastic local operations and classical communication. Unlike a GHZ state, a W state retains bipartite entanglement after one qubit is lost or traced out [1094, 1095]:

$$|W\rangle = \frac{1}{\sqrt{3}}(|001\rangle + |010\rangle + |100\rangle)$$

Entanglement depends on whether the full state factorizes across a chosen partition and, quantitatively, on measures such as entanglement entropy, concurrence, negativity or Schmidt coefficients. Hamming distance can be relevant to the structure of specific states and codes, but states with a large or small Hamming distance can both be separable or entangled.

4.1.7 Matrices

Various matrix transformations must be understood here:

- **Matrix conjugate** ($\bar{A}$) when all complex number see their imaginary part negated, or $a_{ij} \rightarrow a_{ij}^*$.
- **Matrix transpose** (A^T) when all matrix a_{ij} values are transformed into a_{ji} value, with i =line and j =column indices of matrix “cells”.
- **Conjugate transpose** ($A^\dagger$), also named Hermitian adjoint or transconjugate, which is the conjugate of the transpose or vice-versa. It is notated as $A^\dagger$, for A “dagger”.
- **Matrix trace** ($tr(A)$) is the sum of the diagonal entries. It is also the sum of the eigenvalues counted with algebraic multiplicity. A general matrix is not normalized to trace 1. Density matrices are a special case for which $tr(\rho) = 1$.
- **Matrix norm** is $\|A\|_F = \sqrt{\sum_{i,j} |a_{ij}|^2}$ for the Frobenius norm and $\|A\|_2 = \sqrt{\lambda_{\max}(A^\dagger A)}$ for the spectral norm. The latter equals the largest singular value of A .

We also have three important classes of matrices:

- **Hermitian matrices** are equal to their conjugate transpose, meaning that $a_{ij} = a_{ji}^*$.

- **Unitary matrices** are square matrices whose inverse equals their conjugate transpose ($A^\dagger = A^{-1}$). A unitary matrix has several properties, one of which is to have orthogonal eigenvectors and to be diagonalizable. An ideal closed-system quantum gate is represented by a unitary acting on the computational Hilbert space (Figure 175). Implemented gates are noisy quantum channels which may also leak outside that computational subspace.

A unitary operation is the application of a unitary matrix to a computational state vector that we'll later see. Quantum computing reversibility comes from this unitary property. A unitary matrix U can also be expressed as $U = e^{iH}$, with H being a Hermitian matrix, but finding H given U is a complicated calculation problem.

- **Projectors** are Hermitian idempotent operators, $P = P^\dagger = P^2$. Their eigenvalues are 0 or 1. A rank-one projector has one eigenvalue equal to 1, while a rank- r projector has r such eigenvalues. Projectors are generally non-unitary and are central to projective measurement. If $|\psi\rangle$ is normalized, $|\psi\rangle\langle\psi|$ is the rank-one projector onto the span of $|\psi\rangle$.

Notation	Description
z^*	Complex conjugate of the complex number z . $(1+i)^* = 1-i$
$ \psi\rangle$	Vector. Also known as a <i>ket</i> .
$\langle\psi $	Vector dual to $ \psi\rangle$. Also known as a <i>bra</i> .
$\langle\varphi \psi\rangle$	Inner product between the vectors $ \varphi\rangle$ and $ \psi\rangle$.
$ \varphi\rangle \otimes \psi\rangle$	Tensor product of $ \varphi\rangle$ and $ \psi\rangle$.
$ \varphi\rangle \psi\rangle$	Abbreviated notation for tensor product of $ \varphi\rangle$ and $ \psi\rangle$.
A^*	Complex conjugate of the A matrix.
A^T	Transpose of the A matrix.
$A^\dagger$	Hermitian conjugate or adjoint of the A matrix, $A^\dagger = (A^T)^*$.
$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^\dagger$	$\begin{bmatrix} a^* & c^* \\ b^* & d^* \end{bmatrix}$.
$\langle\varphi A \psi\rangle$	Inner product between $ \varphi\rangle$ and $A \psi\rangle$. Equivalently, inner product between $A^\dagger \varphi\rangle$ and $ \psi\rangle$.

Figure 177: Linear algebra key rules. Source: [1096].

Indeed, $(|\psi\rangle\langle\psi|)|\phi\rangle = \langle\psi|\phi\rangle|\psi\rangle$. The inner product $\langle\psi|\phi\rangle$ is generally a complex number. Some of these elements are summarized in Figure 178 and Figure 177.

4.1.8 Pure and mixed states

Let's now explain what the three main states of quantum objects are, basis, pure and mixed as shown in Figure 178. We'll apply it to the case of qubits, given these notions are valid with any quantum system. We are dealing with mathematical models that describe quantum objects states [1097].

Basis states correspond to given combinations of 0 and 1 values in a qubit register. For a single qubit, these are the states $|0\rangle$ and $|1\rangle$. For a register of N qubits, it is one of the 2^N different basis states combinations

of 0s and 1s, or a tensor product of N single qubit basis states. It constitutes the computational basis in a complex numbers Hilbert space of dimension 2^N .

The vectors of this basis are all mathematically orthogonal. A basis state is also named a computational basis state. When measuring individual qubits in these states, you get a deterministic result, at least with theoretically perfect qubits.

Pure states are quantum states represented by a normalized state vector $|\psi\rangle$, or equivalently by a rank-one density operator $\rho = |\psi\rangle\langle\psi|$. A pure state can be a superposition in a chosen basis and, for several subsystems, it can also be entangled. Purity is a property of the state description. An isolated system initially in a pure state remains pure under unitary evolution, while coupling to an unobserved environment can make its reduced state mixed, like when it is affected by ambient noise. Even when $|\psi\rangle$ is exactly known, individual measurements can be intrinsically random unless the state is an eigenstate of the measured observable.

Mixed states are represented by density operators of rank greater than one. They can arise from classical uncertainty over preparations, from tracing out part of an entangled state, or from open-system dynamics. A mixed state can be written as a convex combination $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$, but this ensemble decomposition is generally not unique and the $|\psi_i\rangle$ need not be computational-basis states. Pure and mixed states can both produce probabilistic measurement outcomes. The operational distinction is that a pure state has $\text{tr}(\rho^2) = 1$, whereas a mixed state has $\text{tr}(\rho^2) < 1$.

Typically, mixed states provide the available information describing two sorts of systems:

Random quantum objects, such as photons from an unpolarized source, can be described by a mixed polarization state. A coherent optical element can transform a pure polarization state into another pure state, while an incoherent ensemble of independently prepared $|H\rangle$ and $|V\rangle$ photons with equal probabilities gives $\rho = (|H\rangle\langle H| + |V\rangle\langle V|)/2$. A polarizing beam splitter by itself is a unitary mode splitter. Whether the reduced polarization state is pure or mixed depends on whether path information and relative phase coherence are retained or discarded.

In the other example in Figure 180, two orthogonally polarized optical fields yield a maximally mixed polarization state only if their relative phase is randomized or otherwise inaccessible. If the two fields are mutually coherent with a stable relative phase, their coherent combination instead corresponds to a pure polarization state, generally linear or elliptical.

Subsystems of an inseparable entangled system of several quantum objects. It helps understand what we are measuring at the end of computing when the resulting qubits are still entangled. One case is a set of qubits affected by decoherence coming from interactions with the environment. It helps understand the effect of

$$A = \begin{bmatrix} 2 & i & -2i \\ -i & 1 & 3 \\ 2i & 3 & -1 \end{bmatrix}^\dagger \quad \bar{A} = \begin{bmatrix} 2 & -i & 2i \\ i & 1 & 3 \\ -2i & 3 & -1 \end{bmatrix}$$

transposed matrix

$$A = \overline{(A)}^*$$

Hermitian matrix
matrix equals its transconjugate

$$U|x\rangle = |y\rangle$$

$$|x\rangle = U^\dagger|y\rangle$$

unitary reversibility

Figure 175: Unitary matrices. © Olivier Ezratty, 2021-2025.




Figure 176: Difference between unitary matrices and Hermitian matrices. © Olivier Ezratty, 2021-September 2026.

	basis states	pure states	mixed states
definitions	aka computational basis states, are N dimensions vectors combining 0s and 1s, with 2^N different such vectors for a N qubits register.	vectors in a Hilbert space of norm 1, specified by a single ket describing coherent superpositions of basis states with complex numbers.	or statistical mixture of pure states, are classical statistical ensemble of combination p_i of pure states Ψ_i . Ψ_i can be any combination of pure states but is usually a set of computational basis states.
randomness origin	no randomness with perfect qubits	quantum	quantum and classical
with a single qubit	$ 0\rangle$ and $ 1\rangle$	$ \Psi\rangle = \alpha 0\rangle + \beta 1\rangle$ $ \alpha ^2 + \beta ^2 = 1$	$p_1 \psi_1\rangle, p_2 \psi_2\rangle$ we don't add them, it's just a statistical ensemble, statistical mixture or convex sum of several systems.
with a N qubits register $i = 1$ to 2^N	$ i\rangle$ 01101011> for N=8 all $ i\rangle$ form the computational basis states of the N qubits register, contains N combinations of 0 and 1, all basis states are mathematically orthogonal.	$ \Psi\rangle = \sum_i \alpha_i i\rangle$ $\sum_i \alpha_i ^2 = 1$ α_i = complex number a pure state is a linear superposition of computational basis states.	$\{(p_i \Psi_i)\}$ ensemble notation $\sum_i p_i = 1$ p_i = positive real number probability to find Ψ_i in the mixed state given all p_i are 0 or a 1 in a pure state.

Figure 178: Differences between basis states, pure states and mixed states. © Olivier Ezratty, 2021-2023

decoherence on the state of a qubits register during computing and how error correction codes are mitigating it. Decoherence comes from the entanglement between a system and its environment, thus, the observed system is not yet isolated and becomes a subsystem of a larger entangled system. Thus, it becomes a mixed state (Figure 181). Want to grasp it clearly? You need to toy

with density matrices representations of these pure and mixed states.

Note that these concepts are applicable to both a single qubit and a register of N qubits.




Figure 179: How to generate mixed states with photons. © Olivier Ezratty, 2021-2025.




Figure 180: Another method to generate a mixed state with photons. © Olivier Ezratty, 2021-2025.




Figure 181: Mixed states and pure states when using qubits. © Olivier Ezratty, 2021-2025.

4.1.9 Density matrices

Density matrices, also named density operators, were introduced in 1927 by **John von Neumann** and **Lev Landau** and later expanded by **Felix Bloch**. von Neumann created this formalism to develop his theory of quantum measurements.

A density matrix is a mathematical tool used to describe quantum systems in pure or mixed states. Compared to the state vector that we saw earlier, a density matrix is the only way to mathematically describe a mixed state. It consolidates all the physically significant information

that could be retrieved from a set of quantum objects given what we know about them. Quantum and classical probabilities are boiled in the density matrix [1098].

Usually represented by the sign ρ (rho), a density matrix is a square matrix of complex numbers used to describe a quantum system, like a register of several qubits. Its size is $2^N \times 2^N$ where N is the number of qubits in the register.

The density matrix of a quantum register in **pure state** is the outer product of its computational basis state vector $|\Psi\rangle\langle\Psi|$ as described in Figure 182, with an example using a Bell pair of two qubits. There is no more

information in the density matrix than in the basis state vector at this stage.

A density matrix for a **mixed state** adds several pure states matrices with real probability coefficients p_i . The $|\Psi_i\rangle$ pure states that are combined to form a mixed state can be themselves states from the computational basis (combination of 0s and 1s) but not necessarily.

They can be any vector in the 2^N Hilbert space and made of (normalized) linear superpositions of these basis states. Mathematically speaking, a pure state density matrix is a special case of mixed state density matrix where only one p_i is not zero.

We'll repeat here what was said with pure and mixed states: a mixed state density matrix consolidates both **quantum uncertainties** (that persists even when the system state is well known) and **classical uncertainties** (due to a lack of knowledge of individual quantum sources and preparation conditions) when a pure state density matrix contains only information pertaining to **quantum uncertainties**.

A density matrix has several mathematical properties as described in Figure 183 and detailed afterwards with some differences between pure and mixed states density matrices.

Hermiticity. A density matrix satisfies $\rho = \rho^\dagger$ and is therefore unitarily diagonalizable with real eigenvalues. The eigenvalues are non-negative because a density matrix is also positive semidefinite, not because it is Hermitian. The standard construction $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ makes both Hermiticity and positivity explicit. A global phase of a pure-state ket cancels in the outer product $|\psi\rangle\langle\psi|$.

Positivity. A density matrix ρ is positive semidefinite, meaning $\langle x|\rho|x\rangle \geq 0$ for every vector $|x\rangle$. For a Hermitian matrix this is equivalent to all eigenvalues being non-negative. Its diagonal entries are therefore non-negative in every orthonormal basis. Hermiticity alone would only guarantee that those diagonal entries are real.

Normalization. A density matrix trace equals 1 for both pure and mixed states. A density operator is said to be “normalized to unit trace”. This means that the sum of its diagonal values, which are all positive real numbers, is equal to 1. This normalization arises from two rules: Born’s rule applied to a pure state ($\sum_k |\langle k|\psi\rangle|^2 = 1$) and the sum of classical probabilities for a mixed state ($\sum_j p_j = 1$). As a result, a diagonal element of a density matrix, ρ_{kk} , can be expressed as a weighted sum of probabilities or amplitudes. For a mixed state composed of pure states $|\psi_j\rangle$ with probabilities p_j , the diagonal element ρ_{kk} is given by $\sum_j p_j |\langle k|\psi_j\rangle|^2$. Here, $\langle k|\psi_j\rangle$ represents the amplitude of obtaining outcome k when measuring the pure state $|\psi_j\rangle$. The diagonal elements of the density matrix are called populations in the chosen basis. A statistical mixture refers to the density operator or an

ensemble decomposition, not to an individual diagonal element.

There are some differences between pure and mixed states density matrices.

Projector and purity. A normalized pure-state density matrix is a rank-one projector: $\rho^2 = \rho$, $\text{tr}(\rho) = 1$, and its spectrum is one eigenvalue equal to 1 with all remaining eigenvalues equal to 0, in any Hilbert-space dimension. Mixed states also have $\text{tr}(\rho) = 1$, but their purity $\text{tr}(\rho^2)$ is smaller than 1. For a Hilbert space of dimension d , $1/d \leq \text{tr}(\rho^2) < 1$ for mixed states, with the minimum $1/d$ reached by the maximally mixed state I/d . For N qubits, $d = 2^N$ and the minimum purity is 2^{-N} . The expectation value of an observable A is $\langle\psi|A|\psi\rangle$ for a pure state and $\text{tr}(\rho A)$ for a general state.

Off-diagonal elements can carry phase coherence in a chosen basis and are therefore often called “coherences”. Closed-system unitary evolution follows the Liouville–von Neumann equation, $\dot{\rho} = -(i/\hbar)[H, \rho]$, which preserves purity. Decoherence is open-system evolution and requires a more general master equation. In the common Markovian case this is described by a Gorini–Kossakowski–Sudarshan–Lindblad equation [1099, 1100].

$$\begin{aligned} \frac{|0\rangle\langle 0| + |1\rangle\langle 1|}{2} &= \frac{|+\rangle\langle +| + |-\rangle\langle -|}{2} \\ &= \begin{bmatrix} 1/2 & 0 \\ 0 & 1/2 \end{bmatrix} \\ &= \frac{1}{2}\mathbb{I} \quad (\text{half the identity matrix}) \end{aligned}$$

Purity $\text{tr}(\rho^2)$ quantifies how mixed a state is. It equals 1 for a pure state and reaches its minimum $1/d$ for the maximally mixed state in a d -dimensional Hilbert space. For N qubits, $d = 2^N$, so the minimum is 2^{-N} .

As a result, any time-dependent unitary transformation U applied to this quantum state won’t affect the mixedness. Indeed, the density matrix over time is $\rho(t) = U(t, t_0)\rho(t_0)U^\dagger(t, t_0)$. Its mixedness is $\text{tr}(\rho^2(t)) = \text{tr}(U(t, t_0)\rho(t_0)U^\dagger(t, t_0)U(t, t_0)\rho(t_0)U^\dagger(t, t_0))$ which equals $\text{tr}(\rho^2(t_0))$.

Combinations. A mixed state can admit infinitely many ensemble decompositions. For one qubit, the maximally mixed state $I/2$ can be written as an equal mixture of $|0\rangle$ and $|1\rangle$, or of $|+\rangle$ and $|-\rangle$. For two qubits, the maximally mixed state is $I_4/4$.

Given a density matrix, you can’t compute the pure states that were combined to create it. Said otherwise, quantum states with the same density matrix can’t be distinguished operationally (i.e., by a set of measurements). Also, when a unitary operation U (defined later, sorry) is applied to a mixed state defined by its density matrix ρ , the resulting state density matrix is $U\rho U^\dagger$. For the fun of a better understanding, Figure 184 shows




Figure 182: How a pure state matrix is constructed. © Olivier Ezratty, 2021-2023.




Figure 183: The various mathematical properties of pure and mixed states density matrices. A density matrix or density operator is the most generic way to represent mixed and pure states. It is a linear combination of pure states outer products multiplied by their positive real number classical probability. For a N qubits register, it's a 2^N rows and columns square matrix. Diagonal and off-diagonal terms are populations and coherences in the chosen basis, while purity is basis-independent and given by $\text{tr}(\rho^2)$. * corresponds to the von Neumann entropy level of uncertainty with the state of qubits or subsystem of qubits. © Olivier Ezratty, 2021-2026.

a graphical segmentation of all the various matrix types we've been mentioning in the previous pages and how they are related with each other.

We forgot to define a **non-defective matrix**, which is a diagonalizable matrix. A normal matrix A verifies $AA^\dagger = A^\dagger A$. A matrix that is both Hermitian and unitary is an involutory Hermitian unitary: it is unitarily diagonalizable and its eigenvalues are +1 or -1. "Trivial matrix" is not a standard name for this class.

Single qubit mixed states can be represented by a point inside the Bloch sphere as shown in Figure 185 using a "Death Star" representation, with a statistical mixture of two pure qubit states. The mixed state is a convex sum of pure-state projectors $|\psi_i\rangle\langle\psi_i|$, "convex" meaning that the coefficients p_i are non-negative and

sum to 1. The geometric representation is a good way to figure out why a given mixed state can result from an infinite number of combinations of two pure states. We can combine more than two pure states to create a mixed state. By the way, the Bloch sphere becomes a Bloch ball.

Density matrix dimensionality. For N qubits, let $d = 2^N$. A density matrix is a $d \times d$ complex Hermitian matrix. Hermiticity leaves d^2 independent real parameters, and the unit-trace constraint removes one, so the physical state space has exactly $d^2 - 1 = 4^N - 1 = 2^{2N} - 1$ independent real parameters. Positivity restricts the allowed region but does not further reduce its dimension in the interior.




Figure 184: A Russian dolls map of matrices which embed their mathematical characteristics (normal, Hermitian, unitary, etc.) and their use in quantum physics and quantum computing (for measurement, density matrices, and reversible qubit operations). © Olivier Ezratty, 2021-2026.

We separate the matrix diagonal from the off-diagonal values. The diagonal values are real numbers because they are the positive probability sums of the diagonal values of pure states density matrices, themselves being positive as $|\alpha_i|^2$.

Equivalently, the d diagonal entries contribute d real numbers and the $d(d-1)/2$ independent off-diagonal entries contribute two real numbers each. This gives d^2 real parameters before normalization and $d^2 - 1$ afterwards. It is preferable not to describe this as a number of independent complex parameters, because Hermiticity couples each off-diagonal entry to its complex conjugate (Figure 186).

However, this dimensionality does not correspond to some useful computing resource in standard gate-based programming models although some work has been done to exploit it, but with no additional computing acceleration. [1101] describes a model using not only unitary matrix operator-based quantum gates. It enables the usage of subroutines in programming. But this programming model doesn't seem adopted so far, except for quantum error correction codes which implement measurement during computing. Mixed state based programming is implemented in the qGCL extension of the language pGCL [1102].

An N -qubit normalized pure-state vector contains 2^N complex amplitudes, or 2^{N+1} real coordinates. Normalization removes one real degree of freedom and the physically irrelevant global phase removes another, leaving $2^{N+1} - 2$ independent real parameters for a pure quantum state. If one stores a normalized state vector without quotienting out the global phase, it has

$2^{N+1} - 1$ real coordinates. Thus, wrong is the statement that “A calculation using n number of qubits on a quantum computer would need 2^n classical bits on a standard computer” [1103]. Why? Because one of the 2^N quantum amplitudes in an N -qubit register cannot be stored or emulated on a single bit!




Figure 185: Representation of a single qubit mixed state in the Bloch sphere. A single qubit mixed state can be represented by points inside the Bloch sphere with 3 degrees of freedom: its two usual angles and the vector length. A maximum mixed state is at the center of the Bloch sphere with equiprobability of $|0\rangle$ and $|1\rangle$. A mixed state can result from an infinite number of combinations of various pure states as shown in the sphere. The state purity is measured by its proximity to the Bloch sphere surface. © Olivier Ezratty, 2021-2025.

So why do we care about these density matrices for mixed states? These are mostly used to understand the effects of decoherence and measurement, and with




Figure 186: Computing the dimensionality of a density matrix based on the matrix various constraints. © Olivier Ezratty, 2021.

qubits registers tomography which helps determine their fidelities.

The sequence of quantum gates in a quantum circuit can also be represented by a large unitary matrix of dimension $2^N \times 2^N = 2^{2N}$ complex numbers. So, with a dimensionality close to a density matrix. But this is not an actual computing resource. It deals more with the extensive computing resources required to emulate in-memory an entire unitary algorithm in a classical computer instead of just executing gates one by one on the computational state vector (Figure 187).

There are many other subtleties with density matrices that we can't detail in the book. For example:

Diagonalization is possible for any density matrix, pure or mixed, since a density operator is Hermitian and therefore diagonal in its own eigenbasis. The decomposition gives a classical probabilistic combination of orthogonal pure state eigenvectors, weighted by the eigenvalues. A state that is diagonal in a specified reference basis has no coherence between those basis states, and it may be pure or mixed depending on its eigenvalues. Diagonality alone signals dephasing only relative to a physically selected basis.

Reduced density matrices are the density matrices of subsystems of composite systems. The reduced density matrix for an entangled pure state is a mixed state or mixed ensemble.

Mixed state purification consists, inversely, in integrating a mixed state in a larger system to create or reconstruct a pure state. It is used in some error-correcting codes.

Bipartite pure states are tensor products of two systems that are not entangled. A pure state system is entangled if and only if some of its reduced states are mixed rather than pure. If all were pure, it would mean

that the pure state density matrix ρ would be separable into several pure states, one for each qubit in the case of a qubits register.

Schmidt decompositions are used to decompose bipartite systems and evaluate their level of entanglement. This level of entanglement can be determined with the Schmidt coefficients coming from the Schmidt decomposition.

Matrix rank. The rank of a matrix is the dimension of its column space, equivalently the number of non-zero singular values. For a diagonalizable normal matrix such as a density matrix, it is also the number of non-zero eigenvalues. The rank of a density matrix gives an indication of the purity of the state it represents. A pure state density matrix has a rank 1, since it can be diagonalized into a matrix where only one value in the diagonal is non-zero. A maximally mixed state has a rank of 2^N , i.e., the number of lines and columns in the density matrix representing N qubits. This ordinary operator rank is not an entanglement measure. It must not be confused with the Schmidt rank of a bipartite pure state nor with the Schmidt number of a mixed state, both defined below.

Schmidt rank is defined for a bipartite pure state as the number of non-zero Schmidt coefficients. A product pure state has Schmidt rank 1 and an entangled pure state has Schmidt rank greater than 1. For mixed states, "Schmidt number" is a distinct convex-roof concept. It should not be confused with the ordinary matrix rank of a density operator.

Quantum Channels are completely positive trace-preserving linear maps from operators on an input Hilbert space to operators on an output Hilbert space. They describe open-system dynamics, noise, reset and measurement when outcomes are ignored. For an N -qubit system with $d = 2^N$, a channel can be represented




Figure 187: Dimensionality of a qubit register. The dimensionality of a qubit register density matrix is much larger than its vector state. But it doesn't mean it corresponds to a large information space that has some direct computational utility. The density matrix modelizes the complex interactions between the qubits and their environment, including noise. It also shows the correlations between the many states from the qubit register computational basis. © Olivier Ezratty, 2022.

in an operator basis by a $d^2 \times d^2$ superoperator matrix, but its entries are constrained by complete positivity and trace preservation and are not independent “state dimensions”. Quantum channels are useful to model decoherence, quantum error correction and qubit noise [1104]. It is even possible to build a tomography with a superoperator, aka a quantum process tomography (QPT). One for example can build a QPT of a quantum gate to detect its imperfections. A QPT can also been done for a more complex operation, or unitary applied to a set of qubits, like a Quantum Fourier Transform [1105].

4.1.10 Grad, curls and divs

In the equations of Maxwell, Schrödinger, Dirac, and others that we have seen are used notations good to remember here around the symbol nabla: ∇ , sometimes used with an arrow $\vec{\nabla}$ (Figure 188).

Nabla ∇ is a vector differential operator. Applied to a scalar function of the three dimensions x, y and z of a Euclidean space, it gives the **gradient**, which is a vector. Applied to a vector function, it gives the **divergence**, a scalar, through a dot product, and the **rotational** or **curl**, a vector, through a cross product. The **Laplacian** ∇^2 is a second order operator that applies to scalar and to vector functions. We won't go far in this book with respect to these operators.

4.1.11 Permanent and determinant

This inventory would not be complete without describing an even stranger mathematical object: the **permanent**

of a square **matrix** $n \times n$, invented by Louis Cauchy in 1812. The formula below describes its content.

$$\text{perm}(A) = \sum_{\sigma \in S_n} \prod_{i=1}^n a_{i, \sigma(i)}$$

The Π denotes a multiplication of values from the index matrix i and $\sigma(i)$. σ is a permutation function of integers between 1 and n , the dimension of the matrix (number of columns and rows). The sigma relates to the set of σ functions of the permutation group S_n (also called symmetrical group) which has a size of $n!$ (factorial of n). The values $a_{i, \sigma(i)}$ are the cells of the coordinate matrix i and $\sigma(i)$.

The equations below show what it gives with $n = 2$ and $n = 3$ knowing that beyond that, it becomes less readable.

$$\text{perm} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = ad + bc$$

$$\text{perm} \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} = aei + bfg + cdh + ceg + bdi + afh$$

The permanent is a scalar that can be real or complex depending on the matrix entries. Its defining sum contains $n!$ products, although substantially faster exact classical algorithms exist than evaluating all permutations explic-

scalar field => vector field	scalar field => vector field	vector field => scalar field	vector field => vector field
$\nabla = \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z} \right)$	$\nabla f = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z} \right)$	$\nabla \cdot \vec{G} = \left(\frac{\partial G_x}{\partial x}, \frac{\partial G_y}{\partial y}, \frac{\partial G_z}{\partial z} \right)$	$\nabla \times \vec{G} = \left(\frac{\partial G_z}{\partial y} - \frac{\partial G_y}{\partial z}, \frac{\partial G_x}{\partial z} - \frac{\partial G_z}{\partial x}, \frac{\partial G_y}{\partial x} - \frac{\partial G_x}{\partial y} \right)$
del or nabla : vector differential operator	gradient : scalar field vector of a scalar function space variations	divergence : vector function, showing its local evolution	rotational or curl of a vector function G transforming a vector field in a vector field describing the field variation in space
scalar field => scalar field		vector field => vector field	
$\nabla^2 f = \nabla \cdot \nabla f = \left(\frac{\partial^2 f}{\partial x^2}, \frac{\partial^2 f}{\partial y^2}, \frac{\partial^2 f}{\partial z^2} \right)$		$\nabla^2 \vec{G} = \left(\frac{\partial^2 G_x}{\partial x^2} + \frac{\partial^2 G_x}{\partial y^2} + \frac{\partial^2 G_x}{\partial z^2}, \frac{\partial^2 G_y}{\partial x^2} + \frac{\partial^2 G_y}{\partial y^2} + \frac{\partial^2 G_y}{\partial z^2}, \frac{\partial^2 G_z}{\partial x^2} + \frac{\partial^2 G_z}{\partial y^2} + \frac{\partial^2 G_z}{\partial z^2} \right)$	
scalar function laplacian		vector function laplacian	

Figure 188: del, nabla, gradient, divergence, rotational, curl, Laplacian, with their domains and codomains. You won't need them in the rest of this book, sort of. This is just informative. © Olivier Ezratty, 2023-2025.

itly. Permanents also appear in combinatorics, including counting perfect matchings in bipartite graphs.

Permanents model the output amplitudes and probabilities of ideal boson-sampling experiments. Output amplitudes are proportional to permanents of submatrices of the interferometer unitary, and the output probabilities are their squared moduli with occupation-number normalization factors [1106]. Exact computation of the permanent of a general matrix is $\#P$ -hard [1107]. Ryser's efficient classical algorithm evaluates an $n \times n$ permanent in $O(n2^n)$ arithmetic operations, which is exponentially faster than the naive $n!$ sum but is still exponential [1108]. This contrasts with determinants, which can be computed in polynomial time by Gaussian elimination. The complexity of approximate permanents depends strongly on the matrix class.

The **determinant of a matrix** is a variant of its permanent. $\text{sgn}(\sigma)$ is the sign of permutations, which is +1 if the number of permutations needed to create the permutation is even and -1 if it is odd.

$$\det(A) = \sum_{\sigma \in S_n} \left(\text{sgn}(\sigma) \prod_{i=1}^n a_{i,\sigma(i)} \right)$$

The formula below shows what it gives for $n=3$. Note that the group of permutations includes the permutation that does not change the order of the elements.

$$\det \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} = aei + bfg + cdh - ceg - bdi - afh$$

Determinants have particular properties such as $\det(AB) = \det(A) \cdot \det(B) = \det(B) \cdot \det(A) = \det(BA)$ which can facilitate the calculation of the determinant of a matrix if it can be factorized into several matrices. Also, the determinant of a matrix is the product of its eigenvalues.

So much for the definition of the basics of the linear algebra of quantum computing. I've skipped a lot of other definitions and rules of computation. It was a question of clarifying certain notions that are frequently used in the scientific literature on quantum computing and in many of the reference works cited in this book. What we have just seen may be useful for you to compare some of the scientific literature on quantum computing.

If you like mathematics, linear algebra and complexity, you can have some fun exploring type III factors algebra that describes the observables in relativistic quantum fields theory [1109]! Classical quantum physics and computing is based on simplistic type I factors algebra. Simpler, but still complicated.

4.1.12 Fourier transforms

Fourier transforms are linear operators that decompose functions or discrete vectors into conjugate frequency or momentum representations. They are therefore very much part of linear analysis and, in their discrete form, ordinary finite-dimensional linear algebra. In quantum physics they are especially important because position and momentum, or time and frequency descriptions, are Fourier related.

The Fourier Transform implements a mathematical decomposition of a function $f: x \mapsto f(x)$ into a function $\hat{f}: \xi \mapsto \hat{f}(\xi)$ returning a complex number containing an amplitude and phase for single frequencies ξ . It is a more generic version of Fourier series which works with periodic signals. The Fourier transform can be obtained as an infinite-period limit of a Fourier series under suitable conditions, the two objects having different domains.

It can be used for example to decompose a wave packet pulse signal that is concentrated in time. For a time-domain signal, x is a time in seconds and ξ a frequency in hertz, below. Other conjugate variable pairs, like position and momentum, come with their own units and sign conventions.

$$\hat{f}(\xi) = \int_{-\infty}^{\infty} f(x) e^{-2\pi i x \xi} dx$$

It can be decomposed below using Euler's formula with its real and complex parts separating the amplitude and phase of the Fourier transformed signal:

$$\begin{aligned} \hat{f}(\xi) &= \int_{-\infty}^{\infty} f(x) \cos(2\pi x \xi) dx \\ &\quad - i \int_{-\infty}^{\infty} f(x) \sin(2\pi x \xi) dx \end{aligned}$$

As shown below, the inverse Fourier transforms that frequency decomposition function $\hat{f}(\xi)$ back into its original compound time domain signal $f(x)$.

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(\xi) e^{2\pi i x \xi} d\xi$$

All of this is easier to understand with examples like in Figure 189 decomposing a time domain signal into five frequencies constituents with their respective magnitude and (equal) phases.

Computing Fourier series and transforms is done in many ways:

Discrete-time Fourier Transform (DTFT) is a form of Fourier analysis that is applicable to a sequence of values. It is often used to analyze samples of a continuous function. The term discrete-time refers to the fact that the transform operates on discrete data, often samples whose interval has some units of time.

Discrete Fourier Transform (DFT) converts a finite sequence of equally spaced samples of the function into a same-length sequence of equally spaced samples of the Discrete-Time Fourier transform (DTFT). The samples are complex numbers coming from a DTFT.

Fast Fourier Transform (FFT) computes the discrete Fourier transform (DFT) of a sequence, or its inverse (IDFT). It is an efficient variation of the DFT.

Quantum Fourier Transform (QFT) is the unitary version of the discrete Fourier transform, up to the chosen sign convention. On an n -qubit state it transforms the 2^n complex amplitudes in the computational basis into their discrete Fourier-transform amplitudes. Those amplitudes are not directly readable as a full classical Fourier spectrum because measurement returns samples. We cover it later in the Algorithms part of this book.

Fourier series were created by **Joseph Fourier** (1768–1830, French) as part of his work in the book “The Analytical Theory of Heat” published in 1822. Beforehand, he accompanied Napoleon Bonaparte in his 1798–1801 Egyptian expedition as a scientific advisor. He then became a Prefect for the Isère department, based in Grenoble. Afterwards, he also drove the young Jean-

François Champollion to get interested in deciphering the Rosetta Stone.

4.1.13 Lie groups

Lie groups are mathematical objects that are frequently encountered in quantum physics and quantum computing [1110–1112]. These combine the notions of a group and a smooth manifold. It was created by Sophus Lie in 1873. A typical Lie group is $SU(2)$, which represents all single-qubit unitaries with determinant 1. The full group of single-qubit unitaries is $U(2)$, and every element of $U(2)$ is equivalent to an $SU(2)$ element up to a physically irrelevant global phase. These groups formalism is used in quantum error correction codes, with the Pauli strings used in variational NISQ algorithms [1113], and with optimizing VQE algorithms [1114].

Group: set of elements with a binary operation, often written as multiplication, satisfying four properties: closure, associativity, an identity element, and an inverse for every element. Associativity means $(ab)c = a(bc)$, not that the order of the elements can be exchanged. Exchanging the order, $ab = ba$, is the additional property of commutativity and is not true for general groups.

Smooth Manifold: a manifold is a space that, when viewed locally, resembles Euclidean space (like a flat space) but may have more complicated global topological structure. A smooth manifold is a type of manifold equipped with a smooth structure, meaning it is locally similar to Euclidean space and allows for smooth functions to be defined on it.

4.1.14 Nonlinearities

We often hear about nonlinearities with quantum physics, particularly with the difficulty to implement it with qubits. It is also used in neural networks activation functions in classical computing. But their meaning is not the same in these different scenarios.

Superconducting qubits exploit the nonlinear Josephson potential to make an electrical oscillator anharmonic. The nonlinearity, not damping, makes adjacent transition frequencies unequal, allowing a microwave drive to address the $|0\rangle \leftrightarrow |1\rangle$ transition while suppressing unwanted excitation of $|2\rangle$ and higher levels. For a transmon the anharmonicity is usually negative, so the $|1\rangle \leftrightarrow |2\rangle$ transition frequency is slightly lower than the qubit $|0\rangle \leftrightarrow |1\rangle$ frequency.

Nonlinearities are also sought after in photonics, especially to create quality two-photon quantum gates. Nonlinearities occur when solid media modify the characteristics of photons such as their polarization P and in a nonlinear way with respect to the electric field applied to the solid. The dominant $\chi^{(i)}$ of a nonlinear medium defines its order. A $\chi^{(3)}$ is a third order nonlinear medium.




Figure 189: Fourier transform and inverse Fourier transform applied to signal decomposition and reconstruction. [Source](#), comments © by Olivier Ezratty, 2021.

In a scalar and instantaneous model, $P = \epsilon_0(\chi^{(1)}E + \chi^{(2)}E^2 + \chi^{(3)}E^3 + \dots)$ with ϵ_0 being the vacuum permittivity. Susceptibilities are generally tensors and the real response depends on the interacting frequencies, the polarization directions, the crystal symmetry, dispersion and loss.

The optical Kerr effect is a third-order $\chi^{(3)}$ nonlinearity in which the refractive index changes approximately proportionally to the optical intensity, hence to $|E|^2$. By contrast, the Pockels electro-optic effect is a second-order $\chi^{(2)}$ response in a non-centrosymmetric medium and produces a refractive-index change linear in an applied electric field. This nonlinearity in optics also occurs in many devices such as power lasers.

Finally, nonlinearities are classically used in neural networks activation functions. These are, for example, sigmoid based on exponential fractions.

So how can activation-function-like nonlinearities appear in quantum algorithms when deterministic closed-system quantum evolution is linear and unitary? A deterministic nonlinear quantum gate is not part of standard quantum mechanics. Effective nonlinear input-output behavior can nevertheless arise from measurement, post-selection, dissipation, feedback, data re-encoding, or from classical nonlinear processing in a hybrid algorithm [1115, 1116].

That's what can be done in algorithms solving Navier-Stokes fluid mechanics equations. Nonlinearity can also come from using (irreversible) qubit readout in so-called mid-circuit measurement and feed-forward, where the rest of the circuit is conditioned by the measurement result. This is used in quantum error correction but also in some circuits [1117–1119] and by using multiple copies of the same unitaries [1120].

4.2 Qubits

Qubits are the basic elements of data manipulation in quantum computers. They are the quantum equivalents of classical computing bits. With them, we move from a deterministic to a probabilistic world, with coherent state transformations and correlations that are unavailable to classical bits. The accessible output information remains classical bits, and is bounded by the Holevo theorem [1121].

In conventional digital electronics, logical bits are encoded by two distinguishable physical states, most commonly voltage or charge levels in CMOS circuits. Current flows transiently during switching, so a bit is not generally represented simply by “current flowing” versus “no current”.

The logic is transistors based. A bit readout gives 1 or 0 and deterministically, i.e., if the read operation is repeated several times, or the read operation is repeated after a re-edition of the calculation, it will yield the same result. This is true for data storage of information, for its transport and processing. This is valid, modulo the errors that can occur during this journey. These most often occur in storage and memory and are corrected via error correction systems using some data redundancy, usually with some parity bits for each stored byte, so with a rather low data overhead. In data storage, complicated redundancy systems are used like RAID disks organization mixing and matching several disks and parity error codes to consider the physical errors coming from storage.

In a qubit, everything is different! While qubits are usually initialized at $|0\rangle$, operations on them called quantum gates create a mathematical linear superposition between states $|0\rangle$ and $|1\rangle$. These two states correspond

to two different discrete possible values of a physical property of a quantum object like an electron spin (up or down, in a given direction), a photon polarization or an atom energy level. Qubits are represented mathematically by a vector in a two-dimensional Hilbert space which describes its amplitude and phase, reminding us of the “wave” nature of quantum objects.

We’ll see later how we use the Bloch sphere geometrical representation to understand how amplitude and phase are visualized. It gets more complicated when we apply entanglement multi-qubit gates which, for suitable inputs, create a nonseparable joint state.

At the end of the circuit execution, we read the value of a qubit. A computational-basis projective readout returns 0 or 1 and, in the ideal model, conditionally prepares the corresponding basis state $|0\rangle$ or $|1\rangle$. Other measurement bases and generalized measurements behave differently. The result is probabilistic. Measurement returns limited classical outcomes. The qubit amplitudes are not directly readable and can only be estimated statistically with running the same circuit multiple times.

The role of a quantum algorithm is to transform amplitudes and phases so that measurement statistics reveal useful information. Repeating a circuit over many shots estimates a probability distribution or an expectation value with statistical uncertainty that usually decreases as $1/\sqrt{M}$ for M independent shots. Repetition does not in general turn an intrinsically random outcome into a deterministic one.

It is also dependent on the structure of quantum algorithms which are designed to generate a result with qubits being as close as possible to their so-called “computational basis states”, namely, $|0\rangle$ and $|1\rangle$.

To sort things out, it is still useful to differentiate three levels of ‘qubit objects’ used in computing as described in Figure 190.

4.2.1 Mathematical qubit

Bits and qubits are idealized mathematical objects that implement a pure mathematical formalism with no errors. What is named a “qubit” is above all a mathematical object. Its dimensionality is different than with a bit. It is represented by two complex numbers, the amplitudes α and β from the qubit quantum state description $\alpha|0\rangle + \beta|1\rangle$. Due to normalization ($|\alpha|^2 + |\beta|^2 = 1$) and getting rid of the qubit global phase, its dimensionality becomes two real numbers, usually represented by two angles in the Bloch sphere. Bits and qubit measurement are both mathematical and physical operations. With qubits, it is mathematically based on a projective measurement on the computational basis comprised of $|0\rangle$ and $|1\rangle$, using a Hermitian matrix. Physically, it is using a measurement apparatus operating on the qubit quantum object.

In a classical probabilistic model, a probabilistic pbit would have a probability p of having the value 0 and

1- p of having the value 1 (Figure 191). It would be a linear probabilistic model. Linear probabilistic models are used in the probabilistic processors discussed in the section on unconventional computing.

With qubits, these probabilistic laws are quite different! A qubit vector state is defined by two complex numbers α and β according to the formula describing the qubit quantum object state $|\Psi\rangle$ as $\alpha|0\rangle + \beta|1\rangle$. Quantumly speaking, $|\Psi\rangle$ is a linear superposition of basis states $|0\rangle$ and $|1\rangle$ with coefficients α and β , aka amplitudes. α is a complex number whose square module describes the probability of having the state $|0\rangle$ and β is a complex number whose square module describes the probability of having the state $|1\rangle$.




Figure 191: Bits, probabilistic bits and qubits. 2020-2026.

The sum of the probabilities of the two basis states must give 1. It is indeed not $\alpha + \beta$ but $|\alpha|^2 + |\beta|^2$ that gives 1, the bars around α and β corresponding to their norm. If $\alpha = a + ib$, then its norm is $\sqrt{a^2 + b^2}$.

It comes from the generic probabilistic model developed by **Max Born** in 1926 and from one of the postulates of quantum physics. The Born rule converts state amplitudes and measurement projectors or effects into outcome probabilities. The probability density $|\psi(x)|^2$ of finding an elementary particle in space is only its position-basis special case.

4.2.2 Bloch sphere

The usual geometric representation of the state of a qubit is based on the **Bloch sphere**, which maps a pure qubit ray onto a unit Bloch vector in a sphere of radius 1, with $|\Psi\rangle = \cos(\theta/2)|0\rangle + e^{i\phi} \sin(\theta/2)|1\rangle$ up to a global phase. It enables a representation of a qubit with its amplitude and phase.

Angles. The qubit state $|0\rangle$ is a length 1 vector going from the center of the sphere to the North pole of the sphere and the state $|1\rangle$ is a vector going from the center of the sphere to its South pole. An arbitrary qubit state $|\Psi\rangle$ is represented by a vector with an angle θ (0 to π , latitude) with respect to the vertical z-axis and an angle ϕ (0 to 2π , longitude) with respect to the x-axis located from the center of the sphere to its equator and around

	bits: 0 or 1		qubits: coherent superposition of 0 and 1	
states	mathematical bit	2 possible exclusive states, 0 or 1	mathematical qubit	linear combination of $ 0\rangle$ and $ 1\rangle$
initialization		0 or 1		$ 0\rangle$ basis state
dimensionality		1 binary digit => two possible values		2 real numbers => one point on Bloch sphere
modifications		logic decision tables, irreversible		reversible unitary transformations
readout		0 ou 1, deterministic		0 or 1, probabilistic
errors		no error, perfect mathematical object		no error, perfect mathematical object
physical implementation	physical bit	current/no current in a logic electronic circuit and two states memory device bit	physical qubit	quantum object with two exclusive states for one property
computing vs memory		separate devices and parts of processors		all done in qubits and with quantum gates
computing operations		transistors based logic		amplitude and phase change + entanglement
errors sources		cosmic rays, transistors leakage		decoherence, thermal, electromagnetic, radioactivity...
error levels		2.5×10^{-11} bit per hour error rate		usually $\gg 0,1\%$ with qubit gates and readout (*)
information redundancy	logical bit	parity bits	logical qubit	large number of physical per logical qubits (*)
error correction		ECC (memory), MCA (CPU), RAID (storage)		quantum error correction codes and FTQC
error level after correction		negligible		under an acceptable threshold for fault tolerance

Figure 190: Detailed comparison between classical bits and qubits with separating the mathematical logic, the physical implementation and error correction techniques. (*) have varying values depending on the state of the art. © Olivier Ezratty, 2021-2026.

the z-axis. θ sets the magnitude probabilities through the half-angle functions $\cos^2(\theta/2)$ and $\sin^2(\theta/2)$, and ϕ is the relative phase between the two amplitudes.

Orthogonality. The basis states $|0\rangle$ and $|1\rangle$ are opposite in the Bloch sphere and are mathematically orthogonal. This is highly counter-intuitive and linked to the angle θ that is divided by two in the formulae. When θ equals π , corresponding to a half turn in the sphere, moving from $|0\rangle$ to $|1\rangle$, $\cos(\theta/2) = \cos(90^\circ) = 0$, illustrating that $|0\rangle$ and $|1\rangle$ are mathematically orthogonal states. This is true for any opposing states within the sphere as with the $|\Psi\rangle$ and $|\Psi'\rangle$ examples as shown in Figure 193.

These opposite states are antiparallel or antipodal, meaning parallel but in opposite directions. It explains why angle θ is halved in the equations describing a quantum state in Bloch sphere in the sine and cosine calculations of the formulas giving α and β [1122]!

So, we divide θ by 2 to link the geometric representation in the sphere with the mathematical representation of the qubit state, and to allow a spreading of all the states of a qubit over the whole sphere. The whole sphere occupation of qubits representations makes it easier to describe how single qubit gates work as we'll show later in a graphical way.

For a pure single-qubit state, $\sin \theta$ quantifies the magnitude of the off-diagonal coherence in the computational Z basis: $2|\rho_{01}| = \sin \theta$. It vanishes at $|0\rangle$ and $|1\rangle$ and is maximal on the equator.

Global phase. A qubit representation is usually independent of its global phase. It can be removed from the equation to turn α into a real number. Still, a qubit is

sometimes represented with a global phase of $-\frac{i\varphi}{2}$ as shown in Figure 193. When removing the global phase from α , the complex part of β integrates the phase difference between the amplitudes α and β . In that case, β is a complex number when the qubit is not in the plane crossing the x-axis ($\theta = 0$) and the z-axis ($\phi = 0$) of the Bloch sphere, meaning it has a non-zero phase. The Bloch coordinates themselves are all real expectation values, with $(x, y, z) = (2\text{Re}(\bar{\alpha}\beta), 2\text{Im}(\bar{\alpha}\beta), |\alpha|^2 - |\beta|^2)$. Applying a rotation around the z-axis will generally reintroduce a complex number in the α of the transformed qubit, which we do not necessarily factorize to remove the global phase of the qubit when doing hand calculations.

Information. A pure qubit state has two continuous real parameters, for example the Bloch-sphere angles θ and ϕ , but these parameters are not directly available as classical data from a single physical qubit. A projective measurement gives one binary outcome. More generally, the Holevo bound limits the classical information accessible from an ensemble of quantum states: without prior entanglement, n transmitted qubits cannot convey more than n classical bits of accessible information [1123–1125]. The continuous parameters of an unknown qubit state can instead be estimated from many identically prepared copies using tomography. For simple expectation-value estimation, the statistical uncertainty typically scales as $1/\sqrt{M}$ with the number M of independent shots, hence $M = O(1/\epsilon^2)$ shots for additive precision ϵ . Thus, a qubit is not a high-precision floating-point storage cell.

When the qubit state vector is horizontal in the Bloch sphere, i.e., it sits in its equator, and we have an even

superposition of $|0\rangle$ and $|1\rangle$, but with a variable relative phase between the $|0\rangle$ and $|1\rangle$ amplitudes which is related to the horizontal angle of the vector ϕ with respect to the z axis as in the diagram on the right. Two usually superposed states are $|+\rangle$ and $|-\rangle$. These are orthogonal states. These equatorial states share the same α component of $1/\sqrt{2}$ but opposite β values. This qubit-rich information is then modified by phase rotation quantum gates. If all qubits in the equator share the same 50%/50% amplitude probabilities, they have a different phase (Figure 192).

Generally, most quantum gates do not generate all vector positions in the Bloch sphere. They are often half or quarter turns. The points of the sphere most often used are the cardinal points: the $|0\rangle$, $|1\rangle$, then the four points corresponding to the superposition of $|0\rangle$ and $|1\rangle$ on the equator.




Figure 192: Bloch sphere equator and superposed states
© Olivier Ezratty, 2021-2025.

Clifford-only stabilizer circuits can be efficiently classically simulated according to the Gottesman–Knill theorem. In the stabilizer circuit model, a non-Clifford resource supplied by a T gate, another non-Clifford rotation or magic-state injection is needed both for universality and for escaping that classical simulation. Quantum advantage obtained in other restricted models, such as boson sampling or query problems, must be assessed separately and does not follow the same criterion. Non-Cliffordness is in any case not sufficient by itself to guarantee a quantum speedup, let alone an exponential one. In the case of NISQ systems, we use arbitrary rotation gates but with limited circuit sizes due to uncorrected quantum noise.

A significant part of the quantum computing power comes with playing with the qubit phase that generates interferences between qubits. We'll see that later with algorithms such as phase amplitude and phase kickback.

Origins. The Bloch sphere builds on **Erwin Schrödinger's** wave function of 1926 and on **Max**

Born's associated probabilistic model, created the same year. **Felix Bloch** (1903-1983, Switzerland) represented the state of a two-level quantum system on the sphere in 1946.

Bloch's sphere is frequently assimilated to **Poincaré's sphere**, named after **Henri Poincaré** (1854-1912, France) and created in 1892 [1122, 1126–1128]. It is used to describe the polarization of light (Figure 194). The sphere polar coordinates represent the various types of light polarization: linear polarization (on the sphere equator), left elliptical polarization (upper hemisphere), right elliptical polarization (lower hemisphere) then left and right circular polarization (North and South poles). The vertical axis (circular polarization) and one of the horizontal axes (linear polarization) represent two observables for a photon. All other states can be described as linear superpositions of these couples of basis states. The Bloch and Poincaré spheres share the same two-level and Stokes-parameter geometry, which is a mathematical relationship and not by itself documentary evidence that one inspired the other. The same geometry applies to classical polarization fields and to single-photon polarization qubits, but the latter require quantum states, Born-rule probabilities and quantum measurement. Classical Maxwell equations alone do not describe single-photon measurement statistics.

The Bloch sphere representation is also used for representing an electron spin measured along three orthogonal axes (X, Y, Z), showing how superposition works with spins. Interestingly, the Bloch sphere corresponds to some 3D physics reality for spins but not, for example, for superconducting qubits. In that case, it is a mathematical representation of the qubit phase and amplitude that has no physical 3D meaning.

Sometimes, a system of polar coordinates is used on one circle, positioning the computational basis states of $|0\rangle$ and $|1\rangle$ as geometrically orthogonal vectors. It somewhat duplicates values since $-|0\rangle$ and $-|1\rangle$ are like $|0\rangle$ and $|1\rangle$, with just a different global phase. Only the right half of the circle is useful (Figure 194, right).

Many other fancy qubits representations have been created with projection of the Bloch sphere onto a plane, representations of several qubit states with a single or several Bloch spheres, even some representation of quantum entanglement with three Bloch spheres for two qubits [1129] or with tetrahedrons [1130], torus [1131] and other representations [1132]. None of these have been standardized and have a practical value for quantum developers.

4.2.3 Physical qubit

Bits and qubits are implemented with different sorts of physical devices. With classical bits, we use to say they correspond to currents circulating or not circulating in transistor-based devices. While this is true with processing, it is different with memory and storage. These are




Figure 193: A thorough explanation of the Bloch sphere representation of qubits. © Olivier Ezratty, 2021.




Figure 194: The Poincaré photon sphere, which shares the same two-level geometry as the Bloch sphere, and another, Euclidian, representation of a qubit.

relying on electronic systems storing information like some magnetic encoding in hard disks drives or with two states transistor-based objects in SRAM (used in processors), DRAM (used around processors) or Flash memory (used in SSD and your usual USB memory key).

Qubits are implemented with quantum systems comprised of a single quantum object (atom, electron, photon) or several quantum objects (particularly with superconducting qubits and topological matter qubits like Majorana fermions).

The $|0\rangle$ and $|1\rangle$ states correspond to two distinguishable states of a chosen physical degree of freedom, such as two polarization states, two spin states or two energy levels. Photon polarization can be analyzed with

polarization optics and photodetectors. Electron spin resonance (ESR) is primarily a coherent spin-control and spectroscopy technique, while spin readout uses platform-specific spin-to-charge, optical or dispersive detection mechanisms. Such qubit encodings are often modeled as two-level systems (TLS).

Physical qubits processing is using physical operations: **amplitude and phase changes** implemented by single-qubit gates, which modify amplitudes and relative phases, **entangling-capable multiple-qubit gates** which create entanglement for suitable input states only, **interferences** resulting from the previous operations and are at the core of most quantum algorithms, and **quantum measurement** yielding $|0\rangle$ or $|1\rangle$ for each qubit when computing has ended or, with **mid-circuit measurement**, particularly when executing quantum

error correction codes. Both bits and qubit physical objects are prone to physical errors. While error rates are very small with classical bits, it is currently quite high with qubits.

Two-qubit gate infidelity is usually a dominant hardware error source, but its value varies strongly with platform, device, gate type, connectivity, calibration procedure and benchmark definition. Qubit errors come from their coupling to uncontrolled environmental degrees of freedom and from imperfect control, including thermal excitation, electromagnetic noise, material defects, laser or microwave noise, crosstalk, leakage and, for some solid-state platforms, energetic-particle events such as cosmic rays and environmental radioactivity. Gravity is not a standard practical qubit error source. Isolation requirements are platform-specific and can include vacuum, cryogenic temperature, magnetic shielding, optical filtering and electromagnetic shielding. But we'll see later that for superconducting and electron spin qubits, the required low temperature is also linked to the microwave pulse wavelengths used to control qubits.

These errors require quantum error correction codes, which, as we'll later see, require a significant overhead of physical qubits. In NISQ quantum computers, these errors are also addressed with the quantum error mitigation technique that we'll describe later.

4.2.4 Logical qubit

Quantum error correction is required for scalable, deep fault-tolerant computation. Shorter noisy experiments, NISQ processors and quantum simulators operate without QEC. In classical computing and telecommunications, “bits” are corrected with different techniques including using parity bits.

ECC (error correcting codes) are used in classical memories. Some error correction systems are used in processors like the Intel MCA (Machine Check Architecture) which detects and reports errors in microprocessor. Other systems correct errors in storage like RAID redundancy for hard-disk drives and SSDs. We also have error correction codes used in classical telecoms.

Fault-tolerant architectures encode a logical qubit on multiple physical qubits, with an overhead that depends on the code, physical error rates, connectivity, error syndrome decoder, logical-gate implementation and target algorithmic failure probability. Logical qubits and logical processors have already been demonstrated experimentally [1133–1135]. These demonstrations are important milestones, but they remain far from the scale and logical error rates required for large fault-tolerant applications.

The redundancy overhead becomes much bigger than with parity bits used in classical computing. In logical qubits, physical qubits are processed with quantum error correcting codes. The number of physical qubits assembled into logical qubits depends on their physical

error rate and on the logical qubit error rate that is expected to enable practical quantum computing. For example, the famous integer factoring Shor algorithm is very demanding since using very precise small angles phase rotation gates. It is the same with large quantum chemical simulations.

There is no such thing as a perfect “logical qubit” with no error rate. A logical qubit always has some error rate. But it is low enough to enable a successful computation yielding correct results.

A whole part is dedicated to quantum error detection, mitigation and error correction (QEC), and fault-tolerance (FTQC) in this book, in 5.2.

4.2.5 Qutrits and qudits

While qubits are everywhere in quantum computing, these are not the only quantum objects available to manage quantum information. Quantum computers can also theoretically be built with **qutrits** (with three possible quantum states), **ququarts** (with four possible states) and more generically, with **qudits** (d being the number of possible quantum states of the qubit underlying quantum system).

A d -level qudit carries a Hilbert-space dimension d , equivalent in dimension to $\log_2(d)$ qubits. Therefore, m qudits span dimension $d^m = 2^{m \log_2 d}$, equivalent to $m \log_2(d)$ qubits when d is a power of two. This dimensional saving does not by itself imply a computational speedup because gate sets, error models, connectivity and readout overhead also change [1136].

For example, researchers at Berkeley and at Rigetti are investigating superconducting qudits with more than two levels [1137]. It could also be used with silicon-vacancy spin cavities [1138] and also cold atoms [1139]. Complex trade-offs must be examined to look at the cost/benefits of qudits, between the increased data space and the effects of noise [1140].

The most common qudits are implemented with photons by managing several of their properties. When $d = 8$, a qudit is labeled as a “quoctit” [1141]!

Using qudits would have an impact on quantum algorithms design, programming and code compiling [1142, 1143]. Most quantum algorithms are designed for quantum computers using qubit-based gates. Compilers may translate qubit circuits into a target qudit gate set, subject to the chosen encoding, the native interactions, the connectivity and the error model.

In 2021, a photonic qudit of dimension 20, a quvigint, was efficiently characterized with high-dimensional state tomography [1144]. Any “record” here depends on the encoded degree of freedom, the state type, the control and readout fidelities and the date, and higher-dimensional photonic states exist under other criteria. Qudits could also potentially be used in analog quantum computing [1145].




Figure 195: Qubits, qutrits and ququarts. Source: [1137].

4.3 Registers

In quantum programming, qubits can be grouped into logical registers, much like named collections of bits in classical programs. A quantum processor can host multiple qubit registers. Their boundaries are primarily logical and a circuit can partition available qubits into data, phase, workspace, syndrome or communication registers as needed.

But most other characteristics of a quantum and classical registers are different as shown in Figure 196. The main difference between an n -qubit register and a traditional n -bit register is the way information is represented and transformed, not the amount of data manipulated simultaneously. In conventional computers, 32- or 64-bit registers store integers or floating-point numbers on which elementary mathematical operations are performed.

A register of n qubits holds a vector in a 2^n dimensional space of complex numbers. Its dimensionality is exponentially larger than a n -bits register. Let's take for instance a register of 3 bits and 3 qubits. The first one will store one value at a time as 101 (5 in base 2) while the state of the three qubits register is described by 2 to the power of 3, i.e. 8, complex amplitudes, one per computational basis state. These amplitudes are constrained by normalization and by a physically irrelevant global phase. They describe the state coherently transformed by each operation and are not 8 independent data items that measurement could expose. The total of their modulus squares equals 1 since these are probabilities.

However, these 2^n states amplitudes do not really constitute some information storage capacity. Quantum algorithm's goal here, in the case of a Grover algorithm example, is to amplify the computational basis state amplitude that is the sought result, while reducing all the other amplitudes to near zero (Figure 197).

The output information is a set of n classical bits. The 2^n amplitudes handled during computation are not some useful information that we exploit outside the register. A shot returns classical outcomes, but the result of an algorithm may be a sample distribution, an expectation value, an estimated phase, a decision or a bit string, and many algorithms deliberately output distributions after adaptive or mid-circuit measurements. So, in the end, you don't really process "big data" with quantum computing or at least, you don't output any big data. You may still use some sort of big data to prepare the state of the register before or during calculation. But it is not to the advantage of quantum computing since feeding a quantum register with classical data is quite slow [1146].

The graphic representation in Figure 197 was built using the Quirk open-source simulator to illustrate how a circuit redistributes amplitudes across a 4-qubit computational basis. A QFT performs a deterministic unitary Fourier transform of the input amplitudes.

Another way of presenting things is a little simpler and more graphical. All the register states are on the left. The calculation generates interference between these states to make one of the states on the right come out which is the answer to the problem [1147]. The example is based on the use of only two qubits that give four different "binary" states of the qubits. So, we do not recover 2^n values in practice, but n bits. The operation can be repeated several times to obtain an average in the form of floating point numbers. But it depends on the algorithms. For most of them, a binary output is sufficient, like with the Grover search algorithm which returns an integer index position. Grover amplifies the marked-state probabilities and does not necessarily produce one exact basis state [1148]. The output bit strings must be sampled in full and their joint distribution analyzed, never reconstructed by rounding per-qubit marginals, which discards correlations.

For communication without pre-shared entanglement, the Holevo bound limits the accessible classical information in an ensemble of n transmitted qubits to at most n classical bits. This does not mean that every quantum computation has an n -bit "information content", and entanglement-assisted protocols such as superdense coding require separate accounting of the pre-shared entanglement resource.

At the current stage of qubit development, single and two-qubit gate error rates vary strongly with the platform, the device, the gate type, the connectivity, the calibration procedure, the benchmark definition and the date. They all remain well above what would ideally be needed, either in NISQ regime or for implementing a fault-tolerant system.

By the way, don't believe the nonsense that is the comparison of the exponential size of the qubit registers computational basis state with the number of particles in the Universe. These are not equivalent dimensions. A

n bits register		n qubits register	
2 ⁿ possible states once at a time	n=3 example	2 ⁿ possible states linearly superposed	
evaluable		partially evaluable	
independent copies		no copy	
individually erasable		non individually erasable	
non destructive readout		value changed after readout	
deterministic		probabilistic	
			aka register pure states

Figure 196: Key differences between a classical bit register and a qubit register. Everything is upside down! © Olivier Ezratty, 2021.




Figure 197: Manipulating a 4-qubit register vector state with a quantum circuit created with Quirk. Each shot returns one bit string. In most cases however, particularly for algorithms bringing some exponential speedups, several entries of the computational basis have non-null amplitude. Multiple shots of the circuit must then be run to estimate the output distribution, with a statistical uncertainty and a hardware-dependent SPAM and noise bias. © Olivier Ezratty, 2021.

number of objects *combination* is not homothetic with a number of objects! With a given number of objects, the number of combinations of these objects will always represent a number that is much bigger than the number of objects taken as a reference. And...exponentially!

On the other hand, qubits have constraints with no direct classical analogue. An unknown quantum state cannot be perfectly cloned, and a deterministic closed-system unitary cannot erase an arbitrary state to a fixed pure state. Individual qubits can nevertheless be reset by non-unitary operations involving measurement, dissipation or coupling to a reservoir. Measurement generally disturbs a state unless it is already an eigenstate of the measured observable. These are probabilistic objects that are difficult to manipulate.

Ancilla qubits. Ancillas are auxiliary qubits used as workspace, controls, syndrome qubits, flags, catalysts, magic states or communication resources. Some are uncomputed coherently and returned to a known state, while others are intentionally measured, often in the middle of a circuit, as in quantum error correction. They are therefore neither inherently unread nor a “trash can”. A program can logically partition the available physical qubits into several registers, including data and ancilla registers.

Speed vs space advantage. Quantum computers potential speedup is frequently explained by the data space of registers with their 2^N amplitudes, with saying “you can modify all these value in a single operation instead of 2^N operations in classical computing”. That is not true. Being able with a single gate to modify

Q Changes to the vector state are not instantaneous

You can often read statements like this “*Quantum computers also use something called entanglement. When two qubits are entangled, what happens to one instantly affects the other, even if they’re far apart. This lets quantum computers connect ideas in powerful new ways.*”.

The quoted statement is misleading. A local operation on one member of an entangled pair cannot change any locally measurable probability distribution at the distant member, so entanglement does not enable faster-than-light signaling. Correlations become visible only when measurement records are later compared through an ordinary classical chan-

nel. A separate question is how fast interactions can create or redistribute correlations and entanglement through a many-body processor. For lattice systems with sufficiently local interactions, Lieb-Robinson bounds provide an effective, nonrelativistic light cone with exponentially small tails outside it [1149]. The Lieb-Robinson velocity is a model-dependent propagation bound, not the fundamental speed of light, and the theorem’s assumptions must be reconsidered for long-range or all-to-all interactions. Physical control and communication links still obey relativistic causality. Gate teleportation also requires classical feed-forward and therefore cannot bypass that limit.




Figure 198: Representing qubits manipulations with interferences. Source: William D. Oliver in [1147].

these amplitudes is independent from the number of such steps in your algorithm. As illustrated in Figure 199, there is indeed a difference between the space and time quantum advantages. The space exponential advantage does not warrant a time exponential speedup in your algorithm. It depends on the algorithm depth and the number of quantum gates to execute. Also, contrarily to common wisdom, quantum computing is not instantaneous. Execution time depends on many factors beyond the number of gates like the number of shots, particularly in the NISQ regime, the gate time which can be very long like with trapped ions and the cost of error mitigation and correction. This also leads to various space-time trade-offs optimizations practices [1150].

Quantum parallelism is a shorthand for the fact that a unitary acts coherently on a superposition of computational-basis components. In a classical state-vector simulation, even a local quantum gate can require updating many of the 2^N amplitudes, but the physical gate is not equivalent to executing 2^N independent classical computations whose individual results can all be read out. Quantum speedup comes only when an algorithm uses interference, and often entanglement, so that a desired global property can be extracted with lower query, gate, or total runtime complexity than the best known classical method. The presence of a T gate, a Toffoli gate, or non-Clifford resources can be required for universal computation in particular gate models, but it is neither by itself necessary nor sufficient for an exponential or superpolynomial speedup. Quan-




Figure 199: There's a significant difference between the space and time advantages of quantum computing. The data space of a qubit register does not guarantee a speed up in your algorithm when compared to classical algorithms. A generic pure N qubits state requires 2^N complex amplitudes for an explicit classical state-vector representation, which is a simulation cost and not a directly accessible qubit memory. Quantum algorithms can trade speed at the expense of the number of qubits or the other way around. For example, semiclassical QFT variants can trade coherent ancillas and controlled rotations for sequential measurement and feedforward when the outputs are consumed appropriately. © Olivier Ezratty, 2023-September 2026.

tum parallelism is very different from classical hardware parallelism across processors or cores. The labyrinth metaphor, in which all paths are independently explored and then read out, is misleading because it omits interference and the restrictions imposed by measurement [1151].

4.4 Gates

In classical computing, logic gates execute Boolean algebra using bit-dependent truth tables. Several types of logic gates with one or two inputs are used, including the NAND gate, which is functionally universal: every Boolean function can be built from NAND gates alone. Its physical transistor count depends on the logic family. A standard static-CMOS two-input NAND gate uses four MOS transistors, two pMOS and two nMOS. In practice, integrated circuits use a mixture of gate structures (Figure 200, left).

An Intel Core i5/7 processor with over 10 billion transistors contains several billion logic gates. A processor is obviously very complex, with gates managing access to a cache memory and registers, and instruction pipeline executing the code defining the gates to be used in calculations. These operations are generated at the processor's clock frequency, most often expressed in GHz.

Most conventional Boolean gates such as NAND, NOR, XOR, AND and OR are logically irreversible because their outputs do not uniquely determine their inputs. Landauer's principle associates a minimum thermodynamic cost with logically irreversible information erasure

in a physical implementation, but it does not imply that each abstract Boolean gate dissipates exactly that minimum heat on every execution.

4.4.1 Common gates

Qubits undergo operations via quantum gates that can be applied to one or more qubits. The way the gate are applied in space (on the qubits) and time (in an orderly fashion) is defined in a quantum circuit as presented in the examples in Figure 197 and Figure 199.

Single-qubit gates apply a 2×2 unitary matrices of complex numbers to the qubit state vector containing the famous α and β complex amplitudes. These always generate some rotation of the qubit vector in the Bloch sphere. The norm of the vector remains stable at 1 at least, before any decoherence happens. And quantum gates modify qubits information without reading it. A single qubit gate on a register of N qubits is a unitary operator, a large square matrix of 2^N lines and columns which results from the tensor product of the gate matrix applied to a qubit and the identity operator acting on all the other qubits, in the qubits order.

Two qubit gates apply 4×4 unitary matrices to the computational basis state vector containing 4 entries (2^2).

Three qubit gates apply an 8×8 matrix to a state vector containing $8=2^3$ entries.

We'll now look at the various quantum gates made available to quantum developers. Single qubit gates can be classified in XY and Z gates. XY gates are rotations




Figure 200: Comparison between classical logic gates and qubit gates. © Olivier Ezratty, 2021-2025.

around an axis in Bloch's sphere equator and can be viewed as amplitude change gates while Z gates are rotations around the Z axis and can be described as phase change gates. The variations come from the rotation axis in the Bloch sphere (usually, X, Y or Z) and the angle of the rotation (1/2 turn, 1/4 turn, 1/8 turn or arbitrary rotation angle) [1152].

- **X gate** (or NOT) performs an inversion or bit flip. A $|0\rangle$ becomes $|1\rangle$ and vice-versa. Mathematically, it inverts the α and the β of the two-component vector that represents qubit state. It generates a 180° rotation in the Bloch sphere around the X axis. This gate is often used to initialize to $|1\rangle$ the state of a qubit at the beginning of a process which is by default initialized at $|0\rangle$. It has a derivative, the SX gate used in IBM QPUs, that corresponds to half an X gate. It is a square root of the X gate.

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

- **Y gate** performs a 180° rotation around the Y-axis in the Bloch sphere. It also turns a $|0\rangle$ into $|1\rangle$.

$$Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$$

- **Z gate** applies a sign change to the β component of the qubit vector (phase flip), i.e. a phase inversion and a 180° rotation with respect to the Z axis. The X, Y and Z gates complemented by the identity I are the **Pauli gates**.

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

They have several characteristics like $ZX = iY$ and $X^2 = Y^2 = Z^2 = I$. Their matrices are noted σ_x , σ_y and σ_z . The set $\{I, X, Y, Z\}$ is a basis for all 2×2 complex matrices. A Hermitian single-qubit operator has real Pauli coefficients, while a general unitary has complex coefficients. Equivalently, any single-qubit unitary can be written $e^{i\alpha}(aI - i\mathbf{b} \cdot \boldsymbol{\sigma})$ with real a and $\mathbf{b}$ satisfying $a^2 + \|\mathbf{b}\|^2 = 1$.

- **S gate** generates a phase change, or a quarter turn rotation around the Z-axis (vertical). This is the equivalent of a half Z-gate. It is also called a "phase gate".

$$S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}$$

- **T gate** equivalent to a half S, which generates a phase change of one eighth of a turn. With two of these gates, an S gate is generated. This gate that is not part of Clifford's group (defined ... later) has the particularity of allowing by approximation the creation of any rotation in Bloch's sphere.

$$T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}$$

Adding a non-Clifford resource such as T to Clifford gates gives a universal discrete gate set. The mathematical QFT itself is composed of Hadamard and controlled phase rotations and does not intrinsically require the T gate. In a fault-tolerant Clifford+ T implementation, those phase rotations are synthesized or injected using non-Clifford resources, often dominated by T -state cost.

- **R** rotation gates apply an arbitrary rotation angle in the Bloch sphere around a given axis. They are generated by the Pauli operators, with $R_j(\theta) = e^{-i\theta\sigma_j/2}$. The R_z gate rotates around the z axis, R_x around the x axis and R_y around the y axis [1153]. A $R_z(\text{angle})$ gate is also called a P_{angle} gate (P for phase).

$$R_m = \begin{bmatrix} 1 & 0 \\ 0 & e^{2i\pi/(2^m)} \end{bmatrix}$$

The matrix above is the discrete QFT phase gate $R_m = \text{diag}(1, e^{2\pi i/2^m})$, whose relative phase is $2\pi/2^m$. It must not be confused with the continuous rotation gates $R_x(\theta)$, $R_y(\theta)$ and $R_z(\theta)$, whose angle θ is an explicit parameter. The notation is library-dependent, so an axis label without an angle is not assumed to mean $\pi/2$.

The R_z gates modify the phase of a qubit and not its amplitude. Thus, the measurement of its state $|0\rangle$ or $|1\rangle$ is not affected by this gate. It will return both $|0\rangle$ or $|1\rangle$ with the same proportions, before and after the use of an R_z gate. Only two points of a sphere do not move during a rotation around an axis connecting them.

- **H** gate, *aka* Hadamard-Walsh, turns a qubit at $|0\rangle$ or $|1\rangle$ into an even superposition of $|0\rangle$ and $|1\rangle$. It is fundamental to generate this superposition in the registers in many quantum algorithms. It is equivalent, up to a global phase, to a π rotation about the $(x+z)/\sqrt{2}$ axis. It is often used to initialize a quantum register before executing an oracle-based algorithm like Grover or Simon algorithms. Figure 201 shows a representation of the effect of this gate on a qubit initialized at $|0\rangle$ or $|1\rangle$ and in-between. If we apply two Hadamard gates to a qubit, we return to the starting point. In other words: $HH = I$, with I being the identity operator. In the usual notation, an H gate applied to $|0\rangle$ gives a state $|+\rangle$ and an H gate applied to $|1\rangle$ gives a state $|-\rangle$. Practically, the H gate swaps the Z and X eigenbases. The H gate is also used in error correction stabilizer circuits, in phase kickback and Fourier-like transforms in quantum algorithms, for measurement basis changes, and with qubit state teleportation.

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

- **I** gate is the identity gate. It may be used as a pause. In the real physical world, a real I gate is not an exact identity due to decoherence. During an idle time interval, platform-dependent noise can cause relaxation, dephasing, thermal excitation, coherent drift or leakage.

$$I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

- A reset operation is sometimes indicated at the beginning or middle of a circuit to prepare a qubit in $|0\rangle$. Reset is non-unitary and can be implemented by measurement and conditional feedback, active cooling, optical pumping or engineered dissipation.

The matrix $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = |0\rangle\langle 0|$ is a projector, not a deterministic reset channel. A reset-to- $|0\rangle$ channel can for example be represented by Kraus operators $K_0 = |0\rangle\langle 0|$ and $K_1 = |0\rangle\langle 1|$, giving $\mathcal{E}(\rho) = |0\rangle\langle 0|$ for every input state. The mathematical formalism applied to a single qubit simply illustrates this. But this works only in theory, only if the gate error rate is zero. Since it is not zero, you don't ever have a perfect $|0\rangle$ or $|1\rangle$. A qubit reset operation may also be used to clean up ancilla qubits after their usage, when we are not using the uncompute trick, which is a way to cleanly reset ancilla qubits and remove potential entanglements with other qubits.

Figure 202 shows representations of the effect of these single qubit gates, also labelled unary gates, on qubits initialized in $|0\rangle$ for the gates H, X, Y, R_x and R_y and with $|+\rangle$ for the phase change gates S, T, Z and R_z . Indeed, phase shift gates have no effect on $|0\rangle$ as well as on $|1\rangle$. For $|1\rangle$, it may just change the qubit global phase, and not its relative phase between the qubit amplitudes α and β , with no material impact on most algorithms. In the examples, the R gates use an angle of 90° or $\pi/2$.

We now cover two- and three-qubit gates. Names are standardized here as CNOT or CX for the single-control NOT, CCNOT or Toffoli for the two-control NOT, and CSWAP or Fredkin for the controlled SWAP. In all circuit diagrams, a filled control dot means control-on- $|1\rangle$ and an open dot means control-on- $|0\rangle$, the top wire is the first qubit of the register and the matrices follow a big-endian ordering. Some of these gates are controlled gates such as CNOT, CZ and Toffoli, while others such as iSWAP, $\sqrt{\text{SWAP}}$ and exchange or Ising interactions are not naturally described as controlled operations. A gate is entangling if it can create entanglement for at




Figure 201: Example of application of a Hadamard gate on $|0\rangle$ or $|1\rangle$ qubits. The gate is equivalent, up to a global phase, to a π rotation around the $(x + z)/\sqrt{2}$ axis (in dotted line). © Olivier Ezratty and Felix Berthier (EPITA), 2023-2026.




Figure 202: Bloch sphere representation of various single-qubit gates. T gates are used in FTQC models for the generation of arbitrary rotation gates while these arbitrary rotation gates and quantum error correction are used in NISQ without quantum error correction. © Olivier Ezratty, 2021-2026.

least some product-state inputs. Even an entangling gate need not entangle every input, and subsequent unitary gates can remove previously created entanglement.

- **CNOT gate** is an inversion of the value of a qubit conditioned by the $|1\rangle$ value of another qubit. It is a quantum equivalent of the XOR gate in classical computing. Formerly called Feynman gate (C).
- **CCNOT** or **Toffoli gate** is an inversion of the value of a qubit conditioned by the $|1\rangle$ value of two other qubits. As shown in Figure 203, a Toffoli gate can be constructed using CNOT, H, T and reverse T gates.

- **CZ gate**, or Control-Z, is a conditional phase change Z gate.
- **CS gate**, or Control-S, allows a phase change of a qubit controlled by the state of a qubit.
- **SWAP gate** inverts the quantum values of two qubits. The SWAP gate is the only two-qubit gate that is not creating a new entanglement between the two qubits. If they were separable before the gate, they will still be separable afterwards. It can be generated from the chaining of three consecutive CNOT gates (Figure 208). The key role of SWAP gates is to connect qubits that are physically distant in the register physical layout. A SWAP gate may also




Figure 203: How a Toffoli gate is built using CNOT, T and H gates. A few qubit platforms like trapped ions and cold atoms can implement native Toffoli gates. Source:[1154].

displace some entanglement. For example, if qubits A and B are entangled, but C is not entangled with A and B, a SWAP between B and C will displace entanglement to A and C and leave B unentangled with A and C (Figure 204). SWAP is usually a costly gate. It is not used a lot when the qubit topology enables all to all qubits direct connections like with some trapped ions qubits. Therefore, most SWAP gates are created by compilers.

$$\text{SWAP} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

- **Fredkin gate** is a SWAP gate between two qubits that is conditioned by the state of a third qubit. So, it has three inputs.
- **Generic Control-U gate** is a two qubits gate applying a generic one qubit unitary to a qubit based on the state of a control qubit (Figure 204).
- **Generic multi-controlled quantum gates** are single qubit unitary transformations controlled by an arbitrary number of control qubits. Sort of a C^nU , extending the reach of a Control-U gate. It requires the ability of your compiler to optimize the decomposition of a such a gate into physical gates supported by your target quantum computer [1155, 1156].
- **Phase-controlled R gates** are the equivalent of single-qubit phase-change R gates, conditioned by the state of a control qubit. If the algorithm, like a quantum Fourier transform, requires m to be large, it is not easy to ensure the reliability of the gate because the required precision becomes very large compared to the phase errors generated by the quantum system. However, phase errors are difficult to correct!

A precision record of such a gate seems to have been reached by Honeywell with its trapped ions qubits presented in 2020 which have a rotation precision of $1/500$ turn. This reminds us that during operations, quantum computing is analog. It is digital only at the level of commands and measured results, which become classical bits again [1152, 1157–1159].

There are some reasons to get confused with S, T and R phase gates angles. For example, a S gate is sometimes branded as a $\pi/2$ and sometimes as a $\pi/4$ (Figure 205). The same is applied to a T gate that is sometimes a $\pi/4$ and sometimes a $\pi/8$. The explanation is in the chart below and is related to the way a global phase is applied to the gate unitary operator. We can split hairs using a “rotation” for the large one and a “round” for the small one.

- **Matchgates** is a group of two-qubit parity-preserving gates with a unitary having the following form:

$$G(A, B) = \begin{bmatrix} a & 0 & 0 & b \\ 0 & e & f & 0 \\ 0 & g & h & 0 \\ c & 0 & 0 & d \end{bmatrix}$$

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad B = \begin{bmatrix} e & f \\ g & h \end{bmatrix}$$

$$\det(A) = \det(B) = \pm 1$$

Matchgates are commonly used in quantum chemistry with non-interacting fermions and spin chains, and quantum machine learning. They are efficiently classically simulable when restricted to nearest-neighbor interactions on a 1D chain, which was first discovered by Valiant [1160]. They can still be used to create optimized quantum algorithms [1161, 1162], particularly in quantum chemistry since they correspond to fermionic Gaussian unitaries, meaning they can simulate the evolution of non-interacting fermionic systems. With relaxing the nearest-neighbor constraint slightly (look at what IBM is doing with Nighthawk in the hardware part), then the circuits can become universal for quantum computation and lead to some exponential speedup [1163].

- **Fermionic SWAP gate** is used to support the Pauli exclusion principle and anticommutation relations thanks to the minus sign in its unitary. Among other uses cases, this gate can be used to simulate fermionic systems in quantum chemistry. It can be generated




Figure 204: Effect of the SWAP gate operation and Control-U two-qubit gate unitary matrix on a two qubit vector state. © Olivier Ezratty, 2021-2025.

with combining a SWAP and a CZ gate or 4 CNOT and 2 Hadamard gates.

$$f\text{SWAP} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

The action of a CNOT on superposed inputs is state-dependent. If the control is $\alpha|0\rangle + \beta|1\rangle$ and the target is $|\phi\rangle$, the output is $\alpha|0\rangle|\phi\rangle + \beta|1\rangle X|\phi\rangle$. This is entangled only when both control amplitudes are non-zero and $|\phi\rangle$ is not an eigenstate of X up to phase. For example, a target in $|+\rangle$ is unchanged by X , so CNOT leaves the joint state separable even with a superposed control.

To fully understand the effect of these gates on any qubit state and computational basis vectors for several qubits, you must look at the unitary matrices implementing these gates and their linear effects on the qubits and/or register computational basis vectors.

In other words, the CNOT examples in Figure 206 should be read as specific input cases rather than a universal rule. CNOT creates entanglement only for suitable product-state inputs. One could argue about two things: first, after a couple of operations, we never have a perfect $|0\rangle$ and are rapidly off-bounds, creating tiny entanglement with CNOT gates in that case, and second, most CNOT gates are run after a Hadamard gate was applied on the control qubit, getting off the $|0\rangle$ state!

4.4.2 Physical gates

Other two-qubit gates play a particular role. These are physical gates implemented at the lowest control level depending on the qubit type (Figure 207). They are not necessarily directly useful for developers but are the basis of some specific universal gates sets with some qubit types.

The types of these gates may influence the performance of the quantum computer depending on the overhead it creates for the compiler, transpiler and optimizer [1164, 1165].

- **$\sqrt{\text{SWAP}}$ gate**, or square root SWAP, stops halfway through a SWAP (below). It is a physical level gate used to entangle electron spin qubits.
- **iSWAP gate** is a two-qubit gate that is implemented in some qubit types like those from C12 and Google in one of its Willow chips. There's a LZ gate variant (Landau-Zener) gate proposed for superconducting qubits by German scientists in 2026.

$$\sqrt{\text{SWAP}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{2}(1+i) & \frac{1}{2}(1-i) & 0 \\ 0 & \frac{1}{2}(1-i) & \frac{1}{2}(1+i) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

- **XY gate** is a generic two-qubit gate that implements a rotation by angles β and θ between the states $|01\rangle$ and $|10\rangle$, with the special case $i\text{SWAP} = XY(0, \pi)$. This gate, proposed by Rigetti, can be implemented on superconducting qubits to reduce the number of




Figure 205: Solving the ambiguity of phase gates labelling. © Olivier Ezratty, 2021.




Figure 206: Visualization of a CNOT two-qubit gate effect for selected input states. It doesn't mean that $|0\rangle$ is the only control state for which no entanglement. For example, a target in an X eigenstate can also remain separable. © Olivier Ezratty, 2021.

two-qubit gates required to execute many quantum algorithms [1166, 1167].

$$XY(\beta, \theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\frac{\theta}{2}) & i \sin(\frac{\theta}{2}) e^{i\beta} & 0 \\ 0 & i \sin(\frac{\theta}{2}) e^{-i\beta} & \cos(\frac{\theta}{2}) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

- **ZZ gate** that is implemented with qubits coupling is a technique that can be used with qubit couplers to connect two superconducting qubits and implemented as a CZ gate [1168]. Static ZZ interactions can also generate frequency shifts in fixed frequency transmon qubits [1169].

- **R_{xx} , R_{yy} and R_{zz}** are two-qubit gates that are implemented natively in some trapped ion quantum computers as a Mølmer-Sørensen gate. They are called Ising coupling gates. The gate R_{xx} is implemented natively in IonQ systems. These gates were also created for the first NMR quantum computing systems.

- **Mølmer-Sørensen gate, Cirac-Zoller gate (CNOT), AC Stark shift gate and Bermudez gate** are various two-qubit gates implemented at the physical level with trapped ions qubits. The Mølmer-Sørensen gate is a collective entangling interaction widely used with trapped ions, including same-species ion chains. Mixed-species variants also exist. Its bichromatic implementation can be made relatively insensitive to the motional state compared with sideband gates, provided the phase-space trajectory

closes and residual spin-motion entanglement is suppressed. It is the main entangling gate for many trapped-ion QPUs.

Logical reversibility. Quantum gates have the particularity of being logically reversible. It can easily be visualized for a single qubit gate, which is a simple rotation in the Bloch's sphere and therefore, reversible with the inverse rotation. A multi-qubit gate is a rotation in a wider dimensional space, with 2^N dimensions, N being the number of qubits. Likewise, it is logically reversible with an inverse rotation, but harder to visualize.

We can rewind some parts of algorithms by applying in reverse order the quantum gates that have just been applied to a set of qubits [1170]. One benefit of this process is the so-called uncompute trick used in some oracle-based algorithms. It enables resetting the ancilla qubits used in computation without doing any reading. It avoids damaging the useful qubits that we need to use for the rest of the algorithm.

That being said, qubits can undergo other operations. They could be stored, meaning transferred, in or from quantum memory. They can also be used to encode two bits instead of one, in what is called “superdense coding”, which is mainly used in quantum telecommunications. It related to Quantum Shannon Theory [1171, 1172].

4.4.3 Gate classes and universal gate sets

The science of quantum gates has led to the creation of many concepts and theorems about groups of quantum gates. They are associated with the notion of **universal gate sets**, capable of generating all other quantum gates.

Figure 209 contains a custom diagram summarizing these classes of quantum gates. In short, $U(2^n)$ is the group of all unitary transformations on n qubits. Restricting to determinant-one matrices gives $SU(2^n)$, which represents all such transformations up to a global phase. Likewise, $U(2)$ contains all single-qubit unitaries and $SU(2)$ represents them up to global phase. Clifford's group includes gates with discrete qubits half-turn rotation, the Hadarmard gate, and conditional gates. T (eighth turn) and R as Control-R gates with different angles from π and $\pi/2$ are not in Clifford's group. They are needed to cover $SU(2)$ and $SU(2^n)$ well. In practice, the addition of the T gate is enough to create a universal gate set using approximations.

The classification of the gates begins with the **Pauli gates** that apply half-turn rotations around the X, Y and Z axes of the Bloch sphere of representation of the qubits.

Pauli group includes the gates resulting from the combination of these three Pauli gates and the sign inversion operations on the α or the β of the qubits (± 1 and $\pm i$). On one qubit, the Pauli group includes the gates $\pm I$,

$\pm iI$, $\pm X$, $\pm iX$, $\pm Y$, $\pm iY$, $\pm Z$, and $\pm iZ$, where I is the identity operation.

Clifford group includes single and multiple qubit gates that normalize the Pauli group applicable to n qubits, i.e., the U gates of this group combined with the Pauli group gates σ with $U\sigma U^\dagger$ generate Pauli group gates.

A Clifford gate is a quantum gate that can be decomposed into Clifford group gates. These include Pauli gates (X, Y, Z) and H, S (90° rotation) and CNOT (also called CX for *control-X*) gates. The Clifford group is enumerable and getting very large as n grows. Its size is respectively 24, 11,520 and 92,897,280, 12,128,668,876,800 and 25,410,822,678,459,187,200 elements for $n=1, 2, 3, 4$ and 5 [1173], according to the following formula.

$$|C_n| = \prod_{j=1}^n 2(4^j - 1)4^j = 2^{n^2+2n} \prod_{j=1}^n (4^j - 1)$$

Then, the number of pure quantum states that can be generated with these Clifford unitaries is also enumerable, their count being given by the following formula, which gives 6, 60, 1,080, 36,720 and 2,423,520 pure stabilizer states for $n=1$ to 5 , then 3.2×10^8 at $n=6$, a growth in $2^{n^2/2+O(n)}$ against $2^{2n^2+O(n)}$ for the Clifford group [1174]:

$$|S_n| = 2^n \prod_{j=1}^n (2^j + 1)$$

Then comes the stabilizer polytope, the convex hull of the finite set of pure stabilizer states. It contains uncountably many mixed states because the convex coefficients are continuous, but it is a finite-vertex polytope:

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$$

These states are not prepared deliberately with gates. They come from noise, a magic state degraded during injection or a qubit dephased while idling.

The Clifford group and the stabilizer group scales shouldn't be confused with the size of the computational basis which grows much slower and in powers of 2.

In NISQ mode, we could consider that Clifford group gates are digital while non-Clifford gates are analog when they include arbitrary angle rotation gates. The Clifford group is named through its relation to Clifford algebras, themselves named after the English mathematician William Kingdon Clifford (1845-1879).

Gottesman-Knill's theorem demonstrates that algorithms using gates in the Clifford group can be simulated in polynomial time on classical computers. It means that

	atoms		electrons & spins				photons
qubit type	trapped ions	cold atoms	super-conducting	silicon	NV centers	Majorana fermions	photons
single qubit gates	rotations	$U_{xyz}(\theta, \psi, \mu)$	$R_x(\pm\pi/2), R_z(\lambda)$ (IBM, Rigetti)	R_x, R_y	R_x, R_y	T, H, S (with measurement and QEC)	X, Z, H X, Z, R, CZ (Xanadu)
two qubit gates	XX Möller-Sorensen	CZ C- $U_{xy}(\theta, \psi)$	CZ (IBM, Rigetti), CNOT, iSWAP	$\sqrt{\text{SWAP}}$	CNOT	CNOT (with measurement and QEC)	CNOT (MZI, or with measurement)

Figure 207: Examples of physical qubit gates implement by specific qubit types. Consolidation © Olivier Ezratty, 2021-2024.

they are insufficient to provide an exponential speedup compared to classical computing [1175, 1176]. Another variant of this theorem from Leslie G. Valiant defines conditions for a quantum algorithm to be classically simulable in polynomial time on a classical computer [1177].

Clifford circuits alone are efficiently classically simulable in the standard stabilizer setting, so universal stabilizer-based computation requires a non-Clifford resource. This does not require a three-qubit gate: arbitrary single-qubit unitaries together with any entangling two-qubit interaction are sufficient for universal quantum computation [1178, 1179]. The T gate is a common discrete non-Clifford resource because Clifford+ T is universal, while Toffoli+ H is another universal set using entirely real matrices [1180]. Neither maximal entanglement nor a simple commute/anticommute criterion is required. Noncommutation is necessary for a rich gate algebra but is not a sufficient universality test.

Continuous gates make it possible to generate rotations of arbitrary angle on the Bloch sphere. They are convenient native operations on many NISQ platforms. They can be natively implemented, particularly in NISQ platforms to create the variational algorithms ansatzes that we will cover later in the algorithms part of this book. QFT circuits mathematically require controlled phase rotations of progressively smaller angles. In a fault-tolerant discrete gate set, those rotations are approximated to the required precision rather than implemented as native continuous gates. In fault tolerant quantum computers, these rotations can't be natively implemented and must be constructed with T or Toffoli gates which can be fault tolerant error corrected, although at a very high overhead cost.

Discrete gates are sets of (Hadamard, Z, S, CNOT) that make at best only half and quarter turns in the Bloch sphere.

Universal gate set is a set of gates whose finite products are dense in the desired unitary group, so that any target unitary can be approximated to arbitrary accuracy. A standard sufficient construction is a universal set of single-qubit gates plus an entangling two-qubit gate [1180].




Figure 208: How to create a SWAP gate with three CNOT gates.

Here are some known sets of universal gates:

- CNOT + all single qubit unitaries can enable the creation of any unitary transformation on any number of qubits. This is demonstrated in the **Barenco theorem** according to which $SU(2^n)$ unitaries can be built out of $SU(2)$ unitaries and a CNOT two qubit gate [1179]. It also proves that any unitary transformation $SU(2^n)$ on n qubits can be built with a maximum of 4^n elementary gates.
- CNOT + T (eighth of turn) + Hadamard, using approximations, linked to the **Solovay-Kitaev's theorem**. It proves that a dense and finite set of quantum gates in $SU(2)$ space allows can be used to reconstruct any gate in this space with a maximum error rate ϵ . This theorem was introduced by Robert M. Solovay in 1995 and proven by Alexei Kitaev in 1997 [1181].
- The Solovay-Kitaev theorem gives a gate count polylogarithmic in $1/\epsilon$ under its standard assumptions. The matrices in $SU(2)$ are 2×2 , but $SU(2)$ itself is a three-real-dimensional Lie group. There are various ways to optimize the number of gates to use to create arbitrary rotations [1182].
- It includes unit matrices (from determinant 1) with complex coefficients and dimension 2, i.e., all single qubit gates.

$$SU(2) = \left\{ \begin{pmatrix} \alpha & -\bar{\beta} \\ \beta & \bar{\alpha} \end{pmatrix} : \alpha, \beta \in \mathbb{C}, |\alpha|^2 + |\beta|^2 = 1 \right\}$$

- Toffoli + H gates also form a universal gate set [1180].




Figure 209: Visual taxonomy of qubit gates explaining the Pauli gates, the Pauli group, the Clifford group and the role of T and R gates to create a universal gate set. © Olivier Ezratty, 2021.

This search for a set of discrete quantum gates allowing by approximation to generate a set of continuous gates of arbitrary rotations is important for some algorithms that we will see later, notably the discrete Fourier transform that is exploited in Shor's algorithm.

Figure 210 and Figure 211 illustrates the effect of the sequence of T and H gates which, according to the combinations, allow to cover the different positions of Bloch's sphere, validating **Solovay-Kitaev's** theorem [1182, 1183].

Transpilers are the parts of quantum code compilers that convert any quantum gate in the underlying universal gate set implemented by the quantum processor. They are associated with code optimizers which removes redundancies and reduce the number of physical gates used for computation.

Magic and magic states are needed in fault-tolerant computation based on error correction stabilizer codes. A magic state is a non-stabilizer resource state that can be consumed by state injection to implement a non-Clifford logical operation such as T . Magic-state distillation or cultivation reduces the error of these resource states at substantial qubit and time cost. Magic or another non-Clifford resource is needed to go beyond Clifford computation in this model, but its presence does not enable alone a quantum advantage.

Algorithms such as Shor's factoring algorithm and many QPE-based fault-tolerant algorithms require non-Clifford resources after compilation. Shor's integer factoring algorithm gives a superpolynomial, commonly

described as exponential, improvement over the best known classical factoring algorithms, although no matching classical lower bound is known. QPE itself is a computational primitive and does not generically provide an exponential speedup.

There is a view proposing a separation between quantum states having different levels of entanglement: the entanglement-dominated (ED) phase and the magic-dominated phase (MD). ED states have more entanglement than magic and MD states have more magic than entanglement. The level of magic relates to the "non-Cliffordness" of a circuit. But it shows that low-magic ED states can still be built with T gates [1185].

One last point to visualize here. There are an infinite and non enumerable possible unitary transformations to reach some arbitrary state on the Bloch sphere as shown with the example in Figure 213. This is the same for a whole quantum circuit. There are multiple quantum circuits and algorithms that can generate a given vector state, that is then sampled with running the circuit multiple times. The big question is to find the one with the smaller number of gates considering all the hardware constraints, and QEC/FTQC constraints when running on a future FTQC system.

Teleportation protocol.

One of the most intriguing quantum gate-based quantum algorithms is qubit teleportation. It was introduced by Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres and William K. Wootters in 1993 [1186].



Figure 210: A visual description of Solovay-Kitaev's theorem. Source: [1184].




Figure 211: Visual description of how an arbitrary rotation gate $R(\theta, \phi)$ is built with series of T and H gates, showing how the number of T gates influences the result. The four first circles represent polar coordinates with in the X axis the θ angle and in the Y axis the ϕ angle from the Bloch sphere, represented on the right for the last example. The dots in green are those which are experimentally obtained on a QuEra prototype FTQC QPU in 2025, as compared to ideal results with a + mark. Source: [1135].

It transfers the state of a qubit from one place to another. It doesn't mean that the qubit is physically transmitted. Only the unknown quantum state is transferred, using a pre-shared entangled pair plus two classical bits. The source qubit loses its quantum state in the protocol.

This algorithm is exploiting a pre-existing quantum entanglement channel to transmit the state of a qubit from one end of this channel to the other. Teleportation involves the transmission of two classical bits in the protocol that are used to reconstitute the qubit sent on arrival. As a result, the transmission of the latter cannot be faster than light. Quantum teleportation is one type of feedforward operation.

Due to the quantum no-cloning theorem, this teleportation is a "move" and not a "copy" (or a "cut & paste" instead of a "copy & paste" to use an easy to understand analogy). The state of the transferred qubit is thus destroyed at its origin [1187]. The main use case of this algorithm and its many variants are in quantum cryptography and telecommunications systems that we will discover later. It could also be used in distributed quantum computer architectures. This algorithm can be tested locally in a quantum computer, as proposed by IBM in its Q Systems with Qiskit.

Here is how the quantum teleportation algorithm works:

Step 1: Alice is qubit 1 in the state $\alpha|0\rangle + \beta|1\rangle$. Qubits 2 and 3 are prepared in an entangled Bell state with an Hadamard and a CNOT gate. The 3-qubit system state is the tensor product of qubit 1 and qubit 2+3, qubit 1 being independent and qubits 2 and 3 being linked:

$$|\Psi_0\rangle = (\alpha|0\rangle + \beta|1\rangle) \otimes \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

Step 2: Alice passes qubits 1 and 2 through a CNOT gate and then a Hadamard gate. This mathematically shares the coefficients α and β across all three qubits. The state $|\Psi_1\rangle$ is now a superposition of 8 possible states (2^3). However, we can rewrite this mess specifically to see what Bob (Qubit 3) is holding relative to Alice's qubits:

$$|\Psi_1\rangle = \frac{1}{2} [|00\rangle (\alpha|0\rangle + \beta|1\rangle) + |01\rangle (\alpha|1\rangle + \beta|0\rangle) + |10\rangle (\alpha|0\rangle - \beta|1\rangle) + |11\rangle (\alpha|1\rangle - \beta|0\rangle)]$$

Step 3: Alice measures qubits 1 and 2. This forces them into one of four definite classical states: $|00\rangle$, $|01\rangle$, $|10\rangle$, or $|11\rangle$. Since the three qubits were entangled, Alice's measurement collapses Bob's qubit 3 into a specific corresponding state. The two classical readout bits identify which conditional Pauli correction Bob must apply to recover the input state. They are not physical errors generated by the measurement:

00 :	$\alpha 0\rangle + \beta 1\rangle$	none (identity)
01 :	$\alpha 1\rangle + \beta 0\rangle$	bit-flipped (X)
10 :	$\alpha 0\rangle - \beta 1\rangle$	phase-flipped (Z)
11 :	$\alpha 1\rangle - \beta 0\rangle$	both flipped (XZ)

Step 4: Alice sends the readout value of qubit 1 and 2 to Bob. Bob corrects the value of qubit 3 with a flip




Figure 212: Description of the quantum teleportation protocol. © Olivier Ezratty, 2026.

Three different unitaries reaching the same qubit state

$$U(\hat{n}, \alpha) = e^{-i\frac{\alpha}{2}\hat{n} \cdot \vec{\sigma}} = \cos\frac{\alpha}{2}I - i\sin\frac{\alpha}{2}(\hat{n} \cdot \vec{\sigma}) \quad |\psi_{\text{target}}\rangle = \cos 22.5^\circ |0\rangle + e^{i\varphi} \sin 22.5^\circ |1\rangle$$




Figure 213: Visualizing three different unitary transformations transforming the ground state into an arbitrary state, with the shortest and longest path, and an intermediate path. They are $U(\hat{n}, \alpha) = e^{-i\frac{\alpha}{2}\hat{n} \cdot \vec{\sigma}}$ preparing the same state at $\theta = 45^\circ$, $\varphi = 30^\circ$ from $|0\rangle$. Admissible axes obey $\hat{n} \cdot (\vec{r}_0 - \vec{r}_1) = 0$ and form a continuum, with $\cos \alpha = (\cos \theta - \cos^2 \beta) / \sin^2 \beta$ and path length $L = |\alpha| \sin \beta \geq \theta$ for $\cos \beta = \hat{n} \cdot \vec{r}_0$. The geodesic U_1 ($\alpha = 45^\circ$, $L = 0.79$ rad) saturates the bound, U_2 ($\alpha = 99.9^\circ$) costs 11% more, and U_3 ($\alpha = 307.6^\circ$) winds the long way round for $5.9\times$ the minimum. On real hardware, a native rotation is a calibrated resonant pulse, microwave or laser, driving an equatorial axis whose azimuth is set by the drive phase, while Z rotations are virtual frame changes costing no time at all. Whether the rotation angle is set by the pulse amplitude at fixed duration, as on transmon superconducting qubits where the bandwidth must stay inside the anharmonicity, or by the pulse duration at fixed Rabi frequency, as on trapped ions and spin qubits, determines whether an arbitrary gate costs two fixed-angle pulses or a single pulse proportional to the Bloch arc. NISQ compilers apply an ZYZ Euler decomposition and translate it into that native basis. With FTQC systems, the same decomposition is followed by gate synthesis that discretizes the arbitrary rotations over Clifford+ T gates, with an eye on optimizing T count rather than arc length [1188]. Source: © Olivier Ezratty, plotted with Claude Pro, on August 26th, 2026.

(X) or a phase flip (Z). In the case of a phase flip, he applies the correction:

$$Z(\alpha|0\rangle - \beta|1\rangle) = \alpha|0\rangle + \beta|1\rangle$$

Bob's qubit 3 now contains the initial Alice state $\alpha|0\rangle + \beta|1\rangle$. Alice's original qubit 1 is now a classical "0" or "1". It lost its superposed state.

Gate teleportation is the application of a quantum gate on an unknown state while it is being teleported [1189, 1190]. The practical application of this technique is to enable the creation of two-qubit gates from different systems that are connected through some entanglement. It is an enabling solution for distributed quantum computing. This feedforward operation is particularly suitable for photonic qubits but has also been tested with trapped ions [1191] as well as with superconducting

qubits, as shown in Figure 214. What is important here is to implement such gates in a deterministic way. Since gate teleportation is a costly resource, its use must be minimized by optimizing the partitioning of quantum code across multiple interconnected QPUs [1192]. Gate teleportation is also used in the measurement based quantum computing (MBQC) technique that we will describe later in the photon qubit section in the quantum computing hardware part.

Dynamic quantum circuits (DQC) use mid-circuit measurements and classical feedforward so that later operations are conditionally selected during execution. The circuit program need not be rewritten dynamically: its control flow contains branches conditioned on measurement outcomes.




Figure 215: Example of dynamic quantum circuit usage. Source: [1194] which is about a DQC proposal using an hybrid analog-digital superconducting architecture. Added in 2026.

With the help from gate teleportation, this can enable long range entanglement implementation, using either sacrificial qubits in the qubit chip layout [1195] or with long range microwave links in the case of solid state qubits (Figure 215).

4.4.4 Gate synthesis

Gate synthesis is the art of creating a set of quantum gates to implement an arbitrary unitary transform [1196]. It uses the process of a **unitary decomposition**.

A quantum algorithm is the description of a quantum circuit made of a series of sequenced timely quantum gates operating on 1, 2 and sometimes 3 qubits. It is the way to create a large unitary transformation on the initialized qubits. An arbitrary unitary transformation on N qubits can be implemented with an upper bound of $O(4^N)$ single-qubit and CNOT gates [1197]. This troublesome exponential scale can hopefully be optimized in many ways, avoiding the creation of yet another exponential curse that would annihilate any quantum speedup. Some papers explain how to decompose a unitary on just two qubits [1198, 1199].

An entangling two-qubit gate can create correlations that persist after the physical interaction is switched off. Entanglement is a property of the joint quantum state, not of a continuing physical connection, and it can be undone by later unitary gates. Local decoherence reduces entanglement but does not always remove it completely, and entanglement can also undergo sudden death or be transferred to other subsystems. This is

referred to as being part of many quantum counterfactual models (see its definition in the Glossary). You understand it by toying with the qubit register state vector and computing the mathematics of gate executions “manually”.

4.5 Inputs and outputs

Traditional microprocessors are composed of fixed logic gates, etched into the silicon, and ‘moving’ bits, which are electrical pulses that propagate through the circuit through the various gates. All this at a certain frequency, often in GHz, set by a quartz clock.

In a quantum computer, the first stage of processing resets the quantum register into an initial state. This is called “preparing the system”.

The various registers are first physically configured in the $|0\rangle$ state. The following initialization uses different operators such as the Hadamard transformation to create $(|0\rangle + |1\rangle)/\sqrt{2}$ or the X gate to change this value $|0\rangle$ to $|1\rangle$. Sometimes, more preparation is required to prepare a denser register state, like with quantum machine learning algorithms where some input data must be encoded in the qubit register state vector. Once this initialization is done, computing gates operations are sequentially applied to the qubits according to the algorithm to be executed.

Finally, qubits are measured at the end of the algorithm execution, which has the effect of modifying their quantum state, even when doing what is called a non-demolition measurement, which doesn’t destroy the qubit but still collapses its wave function to one of the computational basis states.

Quantum algorithms diagrams for universal gates computers, in Figure 216 on the right, are most often time diagrams. Both classical and quantum circuit schematics are logical representations that get mapped onto physical interconnects or control schedules. Quantum circuit wires denote logical subsystem histories and not necessarily literal wires. Compilation maps them to platform-specific qubits, couplers, control pulses, ion transport or photonic propagation paths.

Now, let’s toy a little bit with qubits and gates with Quirk, particularly to identify pure and mixed states with single or two qubits in Figure 218. It also shows the role of off-diagonal values in density matrices in Figure 218.

Quirk is an open-source JavaScript based code emulator developed by Craig Gidney that can support up to 16 qubits on your laptop and smartphone (see Figure 217 with the example of a Grover circuit run with 5 qubits) and represent density matrices with up to 8 qubits. It is very easy to use and can help us understand how the structure of a vector state changes by applying quantum gates to a set of qubits.




Figure 214: Gate teleportation with two distant superconducting qubits (D1 and D2), using two communications qubits (C1 and C2) which are entangled, and using two classical communication lines (3 and 4). Source: [1193]. Added in 2023.




Figure 216: Time and space differences with classical logic and quantum gates. © Olivier Ezratty, 2021. Qubits are fixed in space for cold atoms, trapped ions, and solid-state based platforms like superconducting and silicon spin qubits, but not with photons which are flying qubits. They travel in space from a photon source to a destination, a photon detector.

Here, we describe a mixed state generated on two qubits after one of them is entangled with a third qubit.

4.6 Qubit lifecycle

One way to understand how a universal gates quantum computer works is to track the life of a qubit during processing:

Initialization. Many gate-based platforms conventionally reset qubits close to $|0\rangle$, often the ground state, before running a quantum circuit. Initialization can also directly prepare other basis or encoded states, and its physical mechanism and energy cost are platform-dependent. Preparation is never perfectly error-free [1200].

Preparation. It is then programmatically prepared with quantum gates to adjust its values that are vectors in the Bloch sphere. The Hadamard gate is one of the most common one and creates a superposed state of $|0\rangle$ and $|1\rangle$. Single qubit gates apply a rotation of the qubit vector in the Bloch sphere. These rotations are based on unitaries, 2×2 complex number matrix operations applied to the qubit vector $[\alpha, \beta]$. These unitaries satisfy $U^\dagger U = I$, which preserves the state-vector norm. Their trace is not generally equal to 1.

For most quantum algorithms, qubit preparation is usually simple with a set of X gates to set them and H gates to create superposed states. In some cases, like with quantum machine learning, qubit states preparation can be more complex, requiring a lot of gates.




Figure 217: Grover algorithm with 5 qubits running on Quirk on a smartphone, showing the process of amplitude amplification. 2021.




Figure 218: An examples of toying with Quirk to see how pure and mixed states look with two qubits. © Olivier Ezratty, 2021.

Multiple-qubit gates then conditionally link qubits together. Without these quantum gates, little could be done with qubits.

Data manipulation. A separable register can be described qubit by qubit, each qubit carrying two real parameters, the angles θ and φ , or the vector $[\alpha, \beta]$. In an entangled register, individual qubits have no pure state vector of their own and are described by reduced density matrices, while the joint state requires 2^N complex amplitudes representing the proportion of each of the computational basis states made of the various combinations of N 0s and 1s. A normalized N -qubit state vector has $2^{N+1} - 1$ real coordinates after imposing normalization, but one additional degree of freedom is a physically irrelevant global phase. The pure-state manifold therefore has $2^{N+1} - 2$ independent real parameters.

As these gates are operated on the qubits, quantum computing works in an analog way. Indeed, gate-based quantum computers are both digital and analog [120].

Measurement. When we measure the value of a qubit, we obtain a classical binary 0 or 1 with a probabilistic return depending on the qubit state. So, for each qubit, we have a 0 as input, a 0 or a 1 as output, and an infinite number of states in between during calculations.

All this to say that the mathematical richness of qubit-based quantum computing happens only during processing. This is the life cycle of the qubit illustrated in Figure 221.

Another schematic view of how classical and quantum computing are intertwined and the format of data that is handled is provided in Figure 222. What is specific




Figure 219: Another example of toying with Quirk to see how pure and mixed states look with three qubits. © Olivier Ezratty, 2021.

to quantum computing is that the same instructions handle data and computing, i.e., quantum gates. The wealth of data in registers exists only during computing but not at the end, after measurement, where it is back in classical mode, turning the computational basis state vector of dimension $2^{N+1}-1$ real numbers to a meager N classical bit. Also, two deterministic worlds are working next-door with a “probabilistic” gate between the quantum world and classical world.

This is, however, not entirely true. If an algorithm like a Grover search or a Shor integer factoring is designed to output a register state with a non-superposed state and a set of $|0\rangle$ to $|1\rangle$ belonging to the computational basis, we will end up with having a nearly deterministic result. Practically, even with error correction, we will still probably have to run the circuit several times and average the results.

Quantum switch is a curious artefact worth mentioning here. It creates a series of qubit transformations that can be implemented simultaneously in different orders. Like say, A then B and B then A, on a given register state. It defies logic and understanding of time flow, creating an indefinite causal order [1202–1206]. It can even be a useful resource for improving reliability of quantum communications [1207].

4.7 Measurement

We are now investigating quantum measurement, a much broader topic than you may think. We have already explained that quantum measurement is assimilated to a wave function collapse onto basis states, in the case of qubits, $|0\rangle$ or $|1\rangle$. We have also seen that quantum computing is highly probabilistic, requiring executing several times your quantum circuit and making an average of the generated results.

But quantum measurement is way more subtle than that. We will see here what can be measured in qubits and when, what is a projective measurement, what is a POVM, a CPTP map, what are gentle and weak measurements, non-selective and selective measurement, state tomography and the like. Some of these techniques are related to quantum computing, including error corrections and some hardware benchmarking tasks and others with quantum telecommunications.

4.7.1 Projective measurement

A projective measurement is a special ideal measurement defined by a set of orthogonal projectors. It is the most common form of measurement used in quantum computing but not the most general one, since POVMs and quantum instruments describe more general outcome statistics and state updates, both covered further below. We’ll first describe it geometrically and then with some mathematical formalism. Projective measurement is also named a von Neumann measurement since **John von Neumann** elaborated its formalism in 1932.

It is easy to intuitively understand what it looks like using the Bloch sphere for a qubit. A qubit projective measurement corresponds to an orthonormal basis, equivalently to an antipodal direction in the Bloch sphere, and it includes an outcome-dependent state update (Figure 220 and Figure 224). The simplest case of all is a measurement along the z axis containing the $|0\rangle$ and $|1\rangle$ orthogonal vectors. It is about doing a measurement in the qubit computational basis. It could also be, theoretically, a measurement along any other axis, like the $|+\rangle$ and $|-\rangle$ states that sit on the Bloch sphere equator along the x axis. We’ll see later how to achieve this feat.

While ideal closed-system quantum gates are reversible unitary operations, a projective measurement with its classical record is not a unitary transformation on the

measured system. For a qubit it yields one of the eigenvalues of the measured observable with Born-rule probabilities. Repeatability does not come just from the observable operator being self-adjoint. An immediate repetition of the same ideal projective measurement gives the same outcome because the first measurement projects the state into the corresponding eigenspace, provided there is no intervening evolution or disturbance. Of course, the measurement of the qubit modifies its state unless it is already a perfect $|0\rangle$ or $|1\rangle$ in the computational basis.

After an ideal projective measurement on the Z axis, the conditional post-measurement state is $|0\rangle$ or $|1\rangle$. If the input is already one of these eigenstates, the system state is unchanged.

Mathematically, a projective measurement is using Projection-Valued Measures (PVMs) on a closed system. On a given qubit, it uses two orthogonal measurement operators, in the form of 2x2 self-adjointed (Hermitian) matrices.

When measuring a qubit along the Z axis, also named the observable Z with eigenvalues +1 and -1 and eigenvectors $|0\rangle$ and $|1\rangle$ (the observable Z is the matrix representation of a Z single qubit quantum gate!), these PVMs operators are respectively:

$$M_0 = |0\rangle\langle 0| = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

$$M_1 = |1\rangle\langle 1| = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \begin{bmatrix} 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

Given the Z observable operator is $Z=M_0 - M_1$, which returns +1 for $|0\rangle$ and -1 for $|1\rangle$.




Figure 220: Visual difference between a unitary transformation (single qubit gate) and a projective measurement, which makes it easier to understand why a unitary transformation is reversible and a projective measurement is not. © Olivier Ezratty, 2023.

On a general basis, with a quantum object with several distinct states, a measurement operator is a matrix M_m and the probability to get the outcome m (with $m=0$ and 1 in the case of a qubit, or $m=0$ to $N-1$ in the case of an N-state quantum object) is $p(m) = \langle \psi | M_m^\dagger M_m | \psi \rangle$ with the completeness constraint $\sum_m M_m^\dagger M_m = I$ (I being the identity matrix).

For $m = 0$, with $M_0 = |0\rangle\langle 0|$, it reads as:

$$p(0) = \begin{bmatrix} \alpha^* & \beta^* \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = |\alpha|^2.$$

Similarly, $p(1) = |\beta|^2 = 1 - |\alpha|^2$ by Born normalization. A single measurement shot returns only the outcome 0 or 1. The probabilities $p(0)$ and $p(1)$ must be estimated from an ensemble of repeated preparations and measurements, unless they are already known from the state preparation.

Any global phase added to $|\psi\rangle$ disappears from measurement probabilities. If $|\psi'\rangle = e^{i\theta}|\psi\rangle$, then

$$p'(m) = \langle \psi' | M_m^\dagger M_m | \psi' \rangle = e^{-i\theta} e^{i\theta} \langle \psi | M_m^\dagger M_m | \psi \rangle = p(m).$$

After the measurement with the operator M_m , the system state $|\psi\rangle$ becomes the projection of $|\psi\rangle$ on M_m divided by the probability of getting state m:

$$\frac{M_m |\psi\rangle}{\sqrt{\langle \psi | M_m^\dagger M_m | \psi \rangle}} \quad \text{aka} \quad \frac{M_m |\psi\rangle}{\sqrt{\langle \psi | M_m | \psi \rangle}}$$

since $M_m^\dagger = M_m$ (self-adjoint matrix) and $M_m M_m = M_m$ (projector matrix)

All these measurement equations are part of the measurement postulate (usually the third) from quantum mechanics postulates.

Figure 223 helps understand the $\langle \phi | A | \psi \rangle$ Dirac notation, where $|\phi\rangle$ and $|\psi\rangle$ are state vectors and A is an operator. You usually read it from the right. The ket on the right is a vertical vector that is multiplied by the middle object that is a square matrix. It creates a similar vertical vector. Then, you multiply it with the bra on the left which is a horizontal vector. It is an inner product. The result is generally a complex number. When the outer bra and ket are the same state, $\langle \psi | A | \psi \rangle$ is guaranteed to be real only if A is Hermitian.

Now, let's be a bit practical. How can we change the measurement basis with qubits, for implementing a mea-




Figure 221: The effect of projective measurement on a single qubit visualized using the Bloch sphere. The result is random unless the qubit state is already in one of the computational basis states, $|0\rangle$ or $|1\rangle$. © Olivier Ezratty, 2021.




Figure 222: Classical and quantum data flow in gate-based quantum computing. During quantum computing, the process is nearly deterministic, modulo the detrimental effects of noise and decoherence. It becomes probabilistic at the interface with the classical world during qubit measurement. © Olivier Ezratty, 2021-2023.

surement along another axis than Z? At least two options are available:

- It may be possible to *physically* implement a measurement on a different basis than the computational basis. This is, for example, the case with polarization-based photon qubits where the polarizer angle can be dynamically and programmatically modified with some electrically controlled optical settings. It looks more difficult to implement for other types of qubits.
- When the *only* supported measurement is a projective measurement in the computational basis $|0\rangle$ and $|1\rangle$, any other projective measurement can be implemented by first applying a unitary transformation to the qubit that creates a rotation in the Bloch sphere equivalent to moving the measurement axis to the Z axis ($|0\rangle$ and $|1\rangle$) (Figure 224). When we say we do an “X” or “Y measurement”, it means that we first apply an H or $HS^\dagger$ single gate rotation (H = Hadamard gate and S = half a Z gate or quarter phase turn) to handle this axis rotation and then, apply a (computational basis) Z-axis measurement. This is what is

regularly done with quantum error correction codes as well as with MBQC (measurement-based quantum computing).

- With QECs (quantum error correction codes), this sort of projective measurement is applied to ancilla qubits, these additional qubits that detect errors in entangled computing qubits. So, when physicists say they are doing a measurement on a basis of two orthogonal vectors, they mean they are applying first a unitary transformation and then a measurement on the computational basis.

4.7.2 Qubits register measurement

So far, we have just elaborated on measurement mathematical underlying tools and dealt with only one qubit. How about measuring a whole qubit register?

A N-qubit register has 2^N possible computational basis states, from $|00 \dots 00\rangle$ to $|11 \dots 11\rangle$. When measuring once a qubit register, you get one of these states, being a combination of N 0s and 1s.

You could stop there and think, that's my result, fine, I am done! Well, no! Since the measurement outcome is probabilistic and prone to errors, you need to run your quantum circuit for a certain number of times and count the number of times you will get each computational basis state (Figure 225). This will depend on what data is generated by your circuit: a value in the computational basis (series of 0s and 1s) or a superposed state vector and on whether you are using a NISQ or FTQC QPU. Quantum algorithms like Grover concentrate the probability on marked computational basis states rather than producing one exact state. NISQ algorithms like VQE can require up to one million shots to compute just one circuit out of many just to get a sufficient output accuracy. Otherwise, if you repeat the circuit and measurement a large number of times, you will end up estimating the bit-string probabilities only. All relative phase information is absent, so no state vector can be reconstructed that way. Full state tomography requires an informationally complete set of measurement settings and reconstructs a density operator ρ , not necessarily a pure vector [1208]. It also requires an exponential number of circuit shots with regards to the number of qubits, losing any quantum computing speedup gain.

You can then run your algorithm several times and analyze the full output bit strings with the estimator that the algorithm defines, which may be a bit-string distribution, a parity or correlation observable, an estimated phase or an expectation value. Averaging each qubit separately and rounding to the nearest 0 or 1 discards the correlations and is not a valid generic decoding procedure. What is "several"? It depends. IBM proposed to run your algorithm a couple of thousand times on its cloud Q Experience platform when it had from 5 to 133 qubits and stated that this number will grow with the number of qubits, we hope linearly. How do you define

the number of circuit runs, or "shots"? It depends on the algorithm, on the expected accuracy of the outcome, and whether you are using a NISQ variational circuit or a circuit for a FTQC system [1209]. The number of shots could be adjusted dynamically with some intermediate output sampling [1210].

One shot of a quantum circuit is **probabilistic** and, as the number of shots grows, the statistical estimators converge to the device output distribution or expectation value. Systematic noise and model bias do not vanish with more shots, so statistical precision and accuracy must not be confused. If an algorithm generates a single computational basis state (i.e., a combination of 0s and 1s without any superposition), then, the output will be deterministic or nearly deterministic considering the effect of noise, given this usually works well only with FTQC systems. Quantum programs can also be designed to be nondeterministic if they contain mid-circuit measurements yielding some intermediate random results and conditioning the rest of the code execution [1211].




Figure 223: Understanding the $\langle \phi | A | \psi \rangle$ Dirac notation. © Olivier Ezratty.

The order of qubit measurements matters in adaptive and semiclassical circuits. In textbook quantum phase estimation, controlled powers of U act from the phase register onto an eigenstate register, then the inverse QFT is applied to the phase register, and finally the phase-register qubits are measured. In iterative or semiclassical QPE, phase qubits can instead be measured sequentially with classical feedforward. Measurement of one subsystem and discarding its outcome is mathematically represented by a partial trace only for the corresponding non-selective reduced state, not for a conditioned measurement outcome.

We also discussed in this section about mid-circuit measurements. These qubit readout operations are used in multiple contexts. One is of course with error correction, as we'll see with the stabilizer formalism in 5.2.6. It can also play a role in quantum algorithms, including introducing non-linearities, enabling conditional code execution, and teleporting quantum states.


measurement is using a collection $\{M_m\}$ of operators acting on the measured system state space $|\psi\rangle$, with probability of m being: $p(m) = \langle\psi|M_m^\dagger M_m|\psi\rangle$

system state after measurement becomes:

$$\frac{M_m |\psi\rangle}{\sqrt{\langle\psi|M_m^\dagger M_m|\psi\rangle}} \quad \text{with:} \quad \sum_m M_m^\dagger M_m = I$$

a measurement is **projective** if all measurement operators or projectors M_m are satisfying $M_m^2 = M_m$, aka « idempotency »

a classical projective measurement is a projection of the qubit vector on the z axis



the z basis is qubit's **computational basis**:

$$M_0 = |0\rangle\langle 0| = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad M_1 = |1\rangle\langle 1| = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

$$\begin{array}{ll} \text{probabilities} & p(0) = |\alpha|^2 \quad p(1) = |\beta|^2 \\ \text{state after measurement} & \frac{\alpha}{|\alpha|} |0\rangle = e^{i\phi} |0\rangle \quad \frac{\beta}{|\beta|} |1\rangle = e^{i\phi} |1\rangle \end{array}$$

removed global phase

when **another basis projection** is required like x or y axis or any axis in the Bloch sphere, gates are applied to the qubit that change the qubit basis. we then measure qubits using the $|0\rangle$ and $|1\rangle$ basis.


for example, if we want to make a qubit measurement on the $|+\rangle$ and $|-\rangle$ basis, we first apply a X rotation on the qubit and then do a measurement in the $|0\rangle$ and $|1\rangle$ basis.

it enables **non destructive measurement** for the initial qubit and is used in most error correction codes that we'll see later.

Figure 224: Another explanation of projective measurement on a different basis and its usage in non-destructive measurement techniques like with error correction codes. © Olivier Ezratty, 2021-2025.


a single qubit measurement is probabilistic

i.e.: a qubit register after a Hadamard gate applied to all qubits is a simple random numbers generator



on a practical basis:

- the circuit is executed many times, depending on the algorithm
- an average of qubits results is computed, producing a real number
- the averaged result is theoretically deterministic
- modulo the error generated by noise and decoherence



x1000 to x8000 shots and even more required in NISC computers, such as with VQE algorithms


Figure 225: Qubit probabilistic measurement and the notion of computing shots. © Olivier Ezratty, 2021-2025.

4.7.3 From computational vector state to full state tomography


What are we measuring? A single computational state, a statistical weight of 0 and 1 or a full vector state? It depends on the algorithm and on the actual technical need of the undertaken measurement. For most algorithms, repeated shots estimate a bit-string distribution, expectation value or another algorithm-specific statistic. Averaging each qubit and rounding is generally not equivalent to finding the most likely bit string. Computational-basis measurements reveal only the probabilities $|\alpha_x|^2$, not the relative phases of the amplitudes. Recovering a full state requires measurements in multiple bases, as in state tomography (Figure 227).

For algorithms debugging with a reasonable number of qubits and for characterizing the quality of a small group of qubits, it may be useful to compute either a histogram of the whole computational state vector or even, a so-called quantum state tomography which will reconstitute the density matrix of the quantum register (Figure 226).


Repeated computational-basis measurements estimate the probability distribution $p_x = |\alpha_x|^2$ over basis states. They do not reconstruct the complex state vector because all relative phase information is absent. Full state reconstruction requires measurements in additional bases. Since the number of runs grows exponentially,



qubits are measured at the end of computation on each qubit computational basis, several times and averaged: in the general case when the algorithm must generate a pure state.




projective measurement on another basis (after X, Y, Z, R_x, R_y or R_z gates): such as with error correcting codes



an histogram with a 2^N probability split of qubits registers computational basis => useful when the algorithm result combines several states, mostly in the middle of an algorithm for debugging purpose.
 at the **end of computing**, we are supposed to have only one bar with a value close to 1, which is easier to measure with a simple measurement method.

this histogram is possible to compute only for a **reasonable number of qubits** because of its exponential quantum computing cost.



a quantum state tomography is a richer visualization with the full system density matrix => used to assess the quality of qubit gates, entanglement and measurement. more complicated to generate (more repeat projections/measurement). tomography is usually possible for a number of qubits ≤ 6 .

Figure 226: From a vector state to a full density matrix, the various ways to measure the state of a qubit register. Compilation © Olivier Ezratty, 2021-2025.

you understand quickly why it won't make sense to use this technique when we use a large number of qubits.

Development tools like IBM Quantum Experience dumps the vector state of your qubits only for helping you learn about how their system works and understand the impact of noise and decoherence.

Reconstructing a generic N -qubit density matrix is much more demanding because it contains $4^N - 1$ independent real parameters. Standard local-Pauli tomography therefore uses an exponentially large number of measurement settings, up to 3^N local basis choices, and the total sample complexity depends additionally on the desired statistical accuracy and reconstruction method. There is no universal 2^{3N} law. Full arbitrary-state tomography is consequently limited to modest N , while structured assumptions such as low rank, matrix-product operators, compressed sensing or classical shadows can characterize larger systems or selected observables with fewer measurements [1212–1214]. The 2005 Innsbruck eight-ion reconstruction was an important historical milestone but not the current record [1215].

Shadow tomography is such a technique, used to efficiently estimate properties of quantum states, enabling the extraction of multiple specific properties of a quantum system without requiring a full reconstruction of the quantum state. Generic full tomography scales exponentially and is limited to modest system sizes, with no universal qubit-count boundary since feasibility depends on the assumptions, the observable set, the target precision, the platform and the classical post-processing. Structured methods can reach larger systems under low-rank, locality, tensor-network, symmetry or target-observable assumptions. The term shadow refers to the creation of an approximate representation of the quantum state

that is sufficient to estimate the desired properties using randomized measurement basis. It is based on some classical post-processing and with a limited number of quantum measurements [1217–1219].

The graphical representation of these density matrices is often used to evaluate the fidelity of 2 or 3-qubit gates in research papers. The example in Figure 229 illustrates this with comparing the theoretical state of a density matrix for 2 and 4 qubits and measurement results. It also helps qualify the quality of qubits entanglement. Various techniques are proposed to speed-up quantum state tomographies and achieve it with better precision.

However, this is a tool for researchers and hardware designers, not for quantum software developers [1220]. The next step is a Quantum Process Tomography which qualifies the quantum channel of a given process, like a series of gates, one gate, or quantum noise and decoherence. It creates an even richer matrix with 2^{2N} columns and rows, representing a linear operator on the system density matrix, *aka* a superoperator.

4.7.4 Non-selective and selective measurements

A non-selective measurement is a quantum measurement for which the outcome is ignored or unavailable in the state description. It is not defined by whether a human has “read” a detector yet. Mathematically, one averages over all outcomes, producing the channel $\rho \mapsto \sum_m M_m \rho M_m^\dagger$. It deals with the information available about the quantum states we are evaluating, not about the measurement physical operations.

After a measurement of a qubit in a given basis whose outcome is discarded, the coherences of its density matrix in the measurement basis are erased. A non-selective




Figure 227: What happens to your qubits when you progressively measure them. In this example, when all measurements returned 0, we got a |0000> with repeating the operation several times, we get different computational basis states. after thousands runs, we obtain a rough distribution of $|\alpha_i|^2$ probabilities for all i. The intermediate Σ 's and a lot of computing and matrix inversion may enable us to reconstruct a density matrix. © Olivier Ezratty, 2021.




Figure 228: The difference between an ideal 2 and 4-photon density matrices and as measured in experiments. Source: [1216].

projective measurement dephases the state in the measurement basis, $\rho \mapsto \sum_m P_m \rho P_m$. It reduces purity only when the input has coherence across different measurement eigenspaces. For a qubit it becomes maximally mixed only for special inputs, such as $|+\rangle$ under an unread Z measurement. An input already in a Z eigenstate remains pure.

This could be used in a new updated Schrodinger' cat thought experiment, replacing the disintegrating radium

atom by a simple qubit in a superposed state (after a H gate). A measurement at time T would trigger the poison release if the result is |1>. All this in a closed box. Keeping the box closed at time T+whatever would be an equivalent of a non-selective measurement, then opening the box at time T+after whatever, would become a classical measurement of an already totally mixed state.

is described by a completely positive, trace-preserving (CPTP) linear map.

For open quantum systems, however, the interaction between the system and environment leads to non-unitary evolution of the system (e.g., dissipation), which requires CPTP maps for full characterization [1224].

A CPTP map has a Kraus representation $\mathcal{E}(\rho) = \sum_k M_k \rho M_k^\dagger$ with $\sum_k M_k^\dagger M_k = I$. For a projective measurement whose outcome is ignored, $M_k = P_k$ and the state becomes $\rho' = \sum_k P_k \rho P_k$. If the outcome k is retained, the conditional state is $\rho_k = M_k \rho M_k^\dagger / p_k$ with $p_k = \text{tr}(M_k \rho M_k^\dagger)$. A d -dimensional channel can be represented by a $d^2 \times d^2$ superoperator matrix.

Quantum Non-Demolition measurement. A QND measurement is designed so that the measured observable can, ideally, be measured repeatedly without the measurement back-action or subsequent free evolution changing that observable's value over the relevant interval. This does not mean that the entire quantum state is frozen or that conjugate observables are undisturbed. A common sufficient condition is that the measured observable commute with the system Hamiltonian during the relevant evolution, together with a measurement interaction that couples to that observable. QND measurement is reviewed in [1225]. It can be implemented with photons, particularly to measure a photon number, meaning the mode occupation of a single-mode Fock state, using a secondary probe field interfering with the signal field [1226]. A photon-number QND measurement infers that occupation while ideally preserving it for a subsequent measurement. The detectors use a single atom coupled to an optical cavity. Other methods consist in using the cross-Kerr effect where a measured photon traverses an optical medium and changes its refraction index [1227]. It creates a phase shift for a probe photon traversing the same media, its phase being measured with a Mach-Zehnder interferometer [1228]. It has been experimented to measure an electron spin with an additional ancilla quantum dot next to an operational quantum dot [1229]. It also currently works well with superconducting qubits. What is a “demolition measurement”? It happens when, after retrieving the result, the readout creates a significant back-action on the measured quantum system that it either destroys it (like a classical photon counting device that absorbs the counted photons) or turns it into a state outside the computational basis (such as a different energy level than ground/excited levels for a qubit, this being called leakage).

Quantum Steering describes the ability of one party, by choosing local measurements on an entangled bi-

partite state, to remotely prepare different conditional ensembles for the other party that cannot be explained by a local-hidden-state model. The distant party's unconditional reduced density matrix is unchanged, so steering does not enable controllable faster-than-light signaling. Measurement outcomes remain random.

It is a variation of nonlocal correlations intermediate between Bell nonlocality and quantum entanglement [1230, 1231]. The concept was initially formalized by Erwin Schrödinger in 1935. This phenomenon is mainly used in quantum communication and cryptography, where it enables semi-device-independent QKD when trust is partial or limited between two parties.


Quantum Reservoir Engineering deliberately shapes a system's coupling to engineered dissipative environments in order to stabilize desired quantum states or dynamics, cool modes, generate entanglement or implement autonomous error correction. It is about tightly controlling the qubit coupling with its environment and it is connected to quantum error correction techniques. Reduced energetic footprint, shorter readout time or QND measurement are possible application-dependent outcomes, not defining properties of reservoir engineering. The approach was initially imagined for NMR qubits, leveraging the Nuclear Overhauser effect. Then it was tested with trapped ions, using some coupling between the qubit harmonic oscillator and a reservoir of oscillator with laser radiations [1232]. Quantum Reservoir Engineering must not be confused with Quantum Reservoir Computing which is an entirely different beast. Introduced by Keisuke Fujii and Kohei Nakajima in 2017, it is the quantum equivalent of a similar technique used in classical deep learning: a low-dimensional data input is encoded into a higher-dimensional dynamical system, the reservoir, whose transient dynamics act as a feature map that facilitates the separation of input states, and only a classical readout layer is trained [1233, 1234]. It is particularly useful to analyze time series of complex data structures.

The technique is also branded “quantum bath”, “engineered dissipation”, “autonomous feedback” and “coherent feedback”. It has since been tested with superconducting qubits and is the basis of the cat-qubits from Inria, Alice&Bob and Amazon [1235–1237]. It can also be used to create cat-qubits with trapped-ions [1238].

Algorithmic Cooling, including heat-bath algorithmic cooling, redistributes entropy among qubits and repeatedly refreshes designated reset qubits through contact with a bath so that selected target qubits become more polarized than thermal equilibrium would otherwise allow [1239]. It is primarily an initialization and polarization technique.

4.8 Gate-based quantum computing key takeaways

- Gate-based quantum computing is the main quantum computing paradigm. It relies on qubits and finite sequences of quantum gates acting on individual qubits or on two or more qubits, called “circuits”. Other paradigms include analog quantum simulators (Pasqal, QuEra) and quantum annealers (D-Wave). The boundary of a quantum algorithm depends on the abstraction level. An abstract algorithm specifies the mathematical procedure that generates a quantum circuit, the logical circuit specifies encoded operations and measurements, and a compiled implementation adds hardware-specific gates, scheduling and control. Classical optimization, adaptive feedforward, error mitigation, quantum error-correction gadgets and decoding may be integral to an end-to-end implementation even when they are treated as separate layers in a software stack.
- The mathematical effects of quantum gates on qubits are relatively simple linear algebra operations. A set of qubits lives in a large-dimensional Hilbert space made of vectors, using complex numbers, vectors, and matrices. A qubit register supports a vector with 2^N complex numbers. A quantum gate applied a so-called “unitary” operation on this vector. Mathematically, it’s a matrix multiplied by the register vector, which changes the vector. The Dirac Bra-Ket notation helps manipulate vectors and matrices in that formalism. It is often used to describe quantum states, for the decomposition of quantum algorithms and for quantum measurements.
- A qubit is encoded in a two-level quantum degree of freedom with basis states represented by $|0\rangle$ and $|1\rangle$. A pure qubit state can be a superposition of these two states, $\alpha|0\rangle + \beta|1\rangle$, with complex amplitudes satisfying the normalization constraint $|\alpha|^2 + |\beta|^2 = 1$. It is well represented with a vector on a Bloch sphere surface. Quantum circuits exploit phase control, interference and entanglement, which can be generated by two-qubit gates such as CNOT and CZ. Entanglement is a nonseparable property of the joint quantum state. It is usually created, transformed and removed during the circuit execution using conditional multi-qubit gates.
- A qubit register of N qubits can handle a linear superposition of 2^N basis states corresponding to the qubit computational basis, each associated with a complex number. But surprisingly, this exponential growth in size is not enough to create a potential polynomial or exponential speedup with quantum computing. You need a lot of entanglement and some non-obvious quantum gates like arbitrary angle rotation gates (NISQ) or T gates (FTQC).
- While the computational space grows exponentially with the number of qubits, you don’t really store information in the qubit register per se since retrieving it would require an exponential number of circuit execution and qubit readouts. You feed it with some problem and the circuit output is a relatively simple value of interest made of N classical bits. You must deal with it when designing quantum circuits. A quantum circuit won’t generate a lot of data like an image, a video or a stream of text. There are still problems where some values of interest can be extracted that are just real numbers, like in quantum chemistry where you are looking for an energy value.
- A quantum circuit must usually be executed a large number of times, and its results averaged due to the probabilistic nature of qubit measurement. The number of “shots” depends on the algorithm structure and its distribution output, programming paradigm and type of error correction or error mitigation. Running the circuit multiple times and processing the results is called “sampling” a circuit.
- Qubit measurement can be done in various ways, the main one being a classical projective measurement, if possible, a non-demolition one (QND) that will maintain the qubit in its collapsed state after measurement and not destroy it. Other techniques are used that are useful for qubit quality characterization and for quantum error correction like a quantum state or quantum gate tomography.



4.9 Bibliography and notes

These are the bibliographical references and notes for the **Gate based quantum computing** section.

- [165] [Demonstration of Two-Qubit Algorithms with a Superconducting Quantum Processor](#), by Leo DiCarlo et al., Nature, June 2009 (9 pages) (preprint on [arXiv](#)).
- [1089] [Quantum coding](#), by Benjamin Schumacher, PRA, April 1995 (34 pages).
- [1090] [A Short History of Complex Numbers](#), by Orlando Merino, 2006 (5 pages).
- [1091] [Linear Algebra Review and Reference](#), by Zico Kolter et al., 2015 (26 pages).
- [1092] [Eigenvectors and eigenvalues |Chapter 14, Essence of linear algebra](#), by Grant Sanderson, 2016 (17 minutes).
- [1093] [Gilbert Strang lecture at MIT](#), by Gilbert Strang, 2011.
- [1094] [Three qubits can be entangled in two inequivalent ways](#), by Wolfgang Dür et al., 2000 (12 pages).
- [1095] [Symmetry and Classification of Multipartite Entangled States](#), by Adam Burchardt, arXiv, April 2022 (126 pages).
- [1096] [Quantum Computation and Quantum Information](#), by Michael Nielsen and Isaac Chuang, Cambridge University Press, 2010 (710 pages).
- [1097] [The Many Inconsistencies of the Purity-Mixture Distinction in Standard Quantum Mechanics](#), by Christian de Ronde and César Massri, arXiv, October 2024 (21 pages).
- [1098] [The Quantum Density Matrix and its many uses: From quantum structure to quantum chaos and noisy simulators](#), by Apoorva D. Patel, arXiv, August 2023 (28 pages).
- [1099] [Completely Positive Dynamical Semigroups of N-Level Systems](#), by Vittorio Gorini, Andrzej Kossakowski, and E. C. G. Sudarshan, Journal of Mathematical Physics, May 1976 (5 pages).
- [1100] [On the Generators of Quantum Dynamical Semigroups](#), by Göran Lindblad, Communications in Mathematical Physics, June 1976 (12 pages) .
- [1101] [Quantum Circuits with Mixed States](#), by Dorit Aharonov, Alexei Kitaev, and Noam Nisan, arXiv, June 1998 (20 pages).
- [1102] [Quantum programming with mixed states](#), by Paolo Zuliani, 2005 (14 pages).
- [1103] [Simulating subatomic physics on a quantum computer](#), by Sarah Charley, October 2020.
- [1104] [Quantum Channels](#), by Stéphane Attal, (65 pages).
- [1105] [Quantum Process Tomography of the Quantum Fourier Transform](#), by Yaakov S. Weinstein, Timothy F. Havel, Joseph Emerson, Nicolas Boulant, Marcos Saraceno, Seth Lloyd, and David G. Cory, arXiv, June 2004 (26 pages).
- [1106] [The Computational Complexity of Linear Optics](#), by Scott Aaronson and Alex Arkhipov, Theory of Computing, February 2013 (110 pages) (preprint on [arXiv](#)).
- [1107] [The complexity of computing the permanent](#), by Leslie G. Valiant, Theoretical Computer Science, January 1979 (13 pages).
- [1108] [Combinatorial Mathematics](#), by Herbert John Ryser, The Carus Mathematical Monographs No. 14, Mathematical Association of America, June 1963 (154 pages).
- [1109] [The Role of Type III Factors in Quantum Field Theory](#), by Jakob Yngvason, arXiv, 2004 (15 pages).
- [1110] [Berkeley Lectures on Lie Groups and Quantum Groups](#), by Richard Borcherds et al., December 2022 (384 pages).
- [1111] [Quantum Mechanics on Lie Groups: I. Noncommutative Fourier Transforms](#), by Mathieu Beauvillain, Blagoje Oblak, and Marios Petropoulos, arXiv, June 2026 (45 pages).
- [1112] [Quantum Mechanics on Lie Groups: II. Path Integrals](#), by Mathieu Beauvillain, Blagoje Oblak, and Marios Petropoulos, arXiv, July 2026 (41 pages).
- [1113] [Full classification of Pauli Lie algebras](#), by Gerard Aguilar, Simon Cichy, Jens Eisert, and Lennart Bittel, arXiv, July 2024 (14 pages).
- [1114] [Non-Iterative Disentangled Unitary Coupled-Cluster based on Lie-algebraic structure](#), by Mohammad Haidar, Olivier Adjoua, Siwar Baddredine, Alberto Peruzzo, and Jean-Philip Piquemal, arXiv, January 2025 (20 pages).
- [1115] [Quantum Neural Networks](#), by Sanjay Gupta and R.K.P. Zia, arXiv, 2001 (30 pages).
- [1116] [Quantum Algorithms for Deep Convolutional Neural Network](#), by Iordanis Kerenidis, 2020 (36 pages).
- [1117] [Quantum Neuron: an elementary building block for machine learning on quantum computers](#), by Yudong Cao, Gian Giacomo Guerreschi, and Alán Aspuru-Guzik, arXiv, November 2017 (27 pages).
- [1118] [Probabilistic imaginary-time evolution by using forward and backward real-time evolution with a single ancilla: first-quantized eigensolver of quantum chemistry for ground states](#), by Taichi Kosugi, Yusuke Nishiya, Hirofumi Nishi, and Yu-ichiro Matsushita, arXiv, August 2022.
- [1119] [Preparing Arbitrary Continuous Functions in Quantum Registers With Logarithmic Complexity](#), by Arthur G. Rattew and Bálint Koczor, arXiv, May 2022 (40 pages).
- [1120] [Nonlinear transformations in quantum computation](#), by Zoë Holmes, Nolan J. Coble, Andrew T. Sornborger, and Yiğit Subaşı, Physical Review Research, February 2023 (20 pages).
- [1121] [Bounds for the Quantity of Information Transmitted by a Quantum Communication Channel](#), by Alexander S. Holevo, Problems of Information Transmission, July 1973 (7 pages).
- [1122] [The Bloch Sphere](#), by Ian Glendinning, 2005 (33 slides).
- [1123] [Quantum Mechanics and the hydrogen atom](#), (19 slides).
- [1124] [Quantum probabilities and the Born rule in the intuitionistic interpretation of quantum mechanics](#), by Arkady Bolotin, arXiv, February 2017 (14 pages).
- [1125] [Why is theta/2 used for a Bloch sphere instead of theta?](#), StackExchange on Physics.
- [1126] [Lectures on Quantum Computing](#), by Dan C. Marinescu and Gabriela M. Marinescu, 2003 (274 pages).
- [1127] [The statistical interpretation of quantum mechanics](#), by Max Born, 1954 (12 pages).
- [1128] [The mathematics of quantum mechanics](#), by Martin Laforest, University of Waterloo, 2015 (111 pages).
- [1129] [Two-Qubit Bloch Sphere](#), by Chu-Ryang Wie, MDPI, 2020 (14 pages).
- [1130] [Geometry of Qubits - A picture book](#), by Yosi Avron and Oded Kenneth, Technion, 2018 (20 slides).
- [1131] [Geometric Visualizations of Single and Entangled Qubits](#), by Li-Heng Henry Chang, Shea Roccaforte, Ziyu Xu, and Paul Cadden-Zimansky, arXiv, May 2024 (25 pages).
- [1132] [Visualizing Entanglement in multi-Qubit Systems](#), by Jonas Bley, Eva Rexigel, Alda Arias, Nikolas Longen, Lars Krupp, Maximilian Kiefer-Emmanouilidis, Paul Lukowicz, Anna Donhauser, Stefan Küchenmann, Jochen Kuhn, and Artur Widera, arXiv, February 2024 (22 pages).

- [1133] [Suppressing Quantum Errors by Scaling a Surface Code Logical Qubit](#), by Google Quantum AI, *Nature*, February 2023 (6 pages).
- [1134] [Logical Quantum Processor Based on Reconfigurable Atom Arrays](#), by Dolev Bluvstein, Simon J. Evered, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin, *Nature*, February 2024 (8 pages).
- [1135] [A fault-tolerant neutral-atom architecture for universal quantum computation](#), by Dolev Bluvstein, Mikhail D. Lukin, et al., *Nature*, November 2025 (30 pages).
- [1136] [Ultracold polar molecules as qudits](#), by JM Hutson, 2020.
- [1137] [Quantum Simulations with Superconducting Qubits](#), by Irfan Siddiqi, 2019 (66 slides).
- [1138] [A Computation-Enhanced High-Dimensional Quantum Gate for Silicon-Vacancy Spins](#), by Gang Fan and Fang-Fang Du, arXiv, September 2024 (10 pages).
- [1139] [Robust Control and Entanglement of Qudits in Neutral Atom Arrays](#), by Amir Burshtein, Shachar Fraenkel, Moshe Goldstein, and Ran Finkelstein, arXiv, August 2025 (17 pages).
- [1140] [Noisy Qudit vs Multiple Qubits : Conditions on Gate Efficiency for Enhancing Fidelity](#), by Denis Janković, Jean-Gabriel Hartmann, Mario Ruben, and Paul-Antoine Hervieux, arXiv, May 2025 (13 pages).
- [1141] [The 2T-quoctit: a two-mode bosonic qudit for high energy physics](#), by Doga Murat Kürkcüoğlu, Henry Lamm, Oluwadara Ogunkoya, and Leonardo Pierattelli, arXiv, September 2025 (10 pages).
- [1142] [Qudits for decomposing multiqubit gates and realizing quantum algorithms](#), by Evgeniy O. Kiktenko, Anastasiia S. Nikolaeva, and Aleksey K. Fedorov, arXiv, June 2025 (24 pages).
- [1143] [MQT Qudits: A Software Framework for Mixed-Dimensional Quantum Computing](#), by Kevin Mato, Martin Ringbauer, Lukas Burgholzer, and Robert Wille, arXiv, October 2024 (9 pages).
- [1144] [Robust and Efficient High-dimensional Quantum State Tomography](#), by Markus Rambach, Mahdi Qaryan, Michael Kewming, Christopher Ferrie, Andrew G. White, and Jacqueline Romero, arXiv, March 2021 (6 pages).
- [1145] [Analog quantum simulation of the Lipkin-Meshkov-Glick model in a transmon qudit](#), by Elizabeth Champion, Steve Campbell, Gabriel T. Landi, Machiel S. Blok, et al., arXiv, December 2025 (19 pages).
- [1146] [Quantum Computing: Progress and Prospects](#), by Emily Grumbling and Mark Horowitz, *National Academies of Science*, 2019 (272 pages).
- [1147] [Introduction to Quantum Computing](#), by William Oliver, MIT, December 2019 (21 slides).
- [1148] [A Fast Quantum Mechanical Algorithm for Database Search](#), by Lov K. Grover, *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing (STOC '96)*, May 1996 (8 pages).
- [1149] [Lieb-Robinson Bounds and the Generation of Correlations and Topological Quantum Order](#), by Sergey Bravyi, Matthew B. Hastings, and Frank Verstraete, *Physical Review Letters*, August 2006 (4 pages).
- [1150] [Space-Time Optimisations for Early Fault-Tolerant Quantum Computation](#), by Sanaa Sharma and Prakash Murali, arXiv, November 2025 (13 pages).
- [1151] [What is Quantum Parallelism, Anyhow?](#), by Stefano Markidis, arXiv, May 2024 (18 pages).
- [1152] [Gates, States, and Circuits](#), by Gavin E, July 2021 (82 pages).
- [1153] [The Prelude](#), by Microsoft, 2017.
- [1154] [Quantum algorithms for optimizers](#), by Giacomo Nannicini, arXiv, July 2025 (252 pages).
- [1155] [Efficient Implementation of Multi-Controlled Quantum Gates](#), by Ben Zindorf and Sougato Bose, arXiv, March 2025 (31 pages).
- [1156] [Polylogarithmic-depth controlled-NOT gates without ancilla qubits](#), by Baptiste Claudon, Fabrice Debbasch, Alberto Peruzzo, Jean-Philip Piquemal, et al., arXiv, June 2024 (9 pages).
- [1157] [Universality of Quantum Gates](#), by Markus Schmassmann, 2007 (22 slides).
- [1158] [An introduction to Quantum Algorithms](#), by Emma Strubell, 2011 (35 pages).
- [1159] [Equivalent Quantum Circuits](#), by Juan Carlos Garcia-Escartin and Pedro Chamorro-Posada, arXiv, October 2011 (12 pages).
- [1160] [Quantum circuits that can be simulated classically in polynomial time](#), by Leslie G. Valiant, *SIAM Journal on Computing*, 2002 📄.
- [1161] [2D Quon Language: Unifying Framework for Cliffords, Matchgates, and Beyond](#), by Byungmin Kang, Chen Zhao, Zhengwei Liu, Xun Gao, and Soonwon Choi, arXiv, May 2025 (40 pages).
- [1162] [Quon Classical Simulation: Unifying Clifford, Matchgates and Entanglement](#), by Zixuan Feng, Zhengwei Liu, Fan Lu, and Ningfeng Wang, arXiv, May 2025 (32 pages).
- [1163] [Exponential quantum speedups in quantum chemistry with linear depth](#), by Oskar Leimkuhler and K. Birgitta Whaley, arXiv, May 2025 (19 pages).
- [1164] [Quantum Hardware Roofline: Evaluating the Impact of Gate Expressivity on Quantum Processor Design](#), by Justin Kalloor, Mathias Weiden, Ed Younis, John Kubiawicz, Bert De Jong, and Costin Iancu, arXiv, February 2024 (13 pages).
- [1165] [YAQQ: Yet Another Quantum Quantizer – Design Space Exploration of Quantum Gate Sets using Novelty Search](#), by Aritra Sarkar, Akash Kundu, Matthew Steinberg, Sibasis Mishra, Sebastiaan Fauquenot, Tamal Acharya, Jarosław A. Miszczak, and Sebastian Feld, arXiv, June 2024 (29 pages).
- [1166] [Implementation of the XY interaction family with calibration of a single pulse](#), by Deanna M. Abrams, Nicolas Didier, Blake R. Johnson, Marcus P. da Silva, and Colm A. Ryan, arXiv, December 2019 (13 pages).
- [1167] [XXPlusYYGate](#), by IBM.
- [1168] [Implementation of Conditional-Phase Gates based on tunable ZZ-Interactions](#), by Michele C. Collodo, Johannes Herrmann, Nathan Lacroix, Christian Kraglund Andersen, Ants Remm, Stefania Lazar, Jean-Claude Besse, Theo Walter, Andreas Wallraff, and Christopher Eichler, arXiv, May 2020 (10 pages).
- [1169] [ZZ-Interaction-Free Single-Qubit-Gate Optimization in Superconducting Qubits](#), by Shu Watanabe, Yutaka Tabuchi, Kentaro Heya, Shuhei Tamate, and Yasunobu Nakamura, arXiv, January 2024 (11 pages).
- [1170] [Synthesis and Optimization of Reversible Circuits - A Survey](#), by Mehdi Saeedi and Igor L. Markov, arXiv, March 2013 (34 pages).
- [1171] [Quantum Shannon Theory](#), by John Preskill, arXiv, July 2025 (97 pages).
- [1172] [Classical to Quantum Shannon Theory](#), by Mark Wilde, *Louisiana State University*, December 2021 (774 pages).
- [1173] [Clifford group](#), by Maris Ozols, 2008 (4 pages).
- [1174] [Hudson's Theorem for finite-dimensional quantum systems](#), by David Gross, *Journal of Mathematical Physics* 47, 122107, arXiv, February 2006 (17 pages).
- [1175] [Positive Wigner functions render classical simulation of quantum computation efficient](#), by A. Mari and J. Eisert, arXiv, September 2012 (7 pages).

- [1176] [Quantum computational advantage implies contextuality](#), by Farid Shahandeh, arXiv, December 2021 (6 pages).
- [1177] [Quantum Computers that can be Simulated Classically in Polynomial Time](#), by Leslie G, 2002 (10 pages).
- [1178] [Universal Quantum Computation and Simulation Using Any Entangling Hamiltonian and Local Unitaries](#), by Jennifer L. Dodd, Michael A. Nielsen, Michael J. Bremner, and Robert T. Thew, Physical Review A, April 2002 (4 pages)  (preprint on [arXiv](#)).
- [1179] [Elementary gates for quantum computation](#), by Adriano Barenco et al., arXiv, 1995 (31 pages).
- [1180] [A Simple Proof that Toffoli and Hadamard are Quantum Universal](#), by Dorit Aharonov, arXiv, January 2003 (4 pages).
- [1181] [Quantum computations: algorithms and error correction](#), by Alexei Kitaev, 1997 (60 pages).
- [1182] [T-count and T-depth of any multi-qubit unitary](#), by Vlad Gheorghiu, Michele Mosca, and Priyanka Mukhopadhyay, arXiv, February 2023 (28 pages).
- [1183] [Shorter quantum circuits via single-qubit gate approximation](#), by Vadym Kliuchnikov, Kristin Lauter, Romy Minko, Adam Paetznick, and Christophe Petit, arXiv, December 2023 (88 pages).
- [1184] [A Guide to Quantum Computers and the Development of Quantum Advantage](#), by Siddharth R. Chander, Terra Science and Education, 2023 (19 pages).
- [1185] [Magic-induced computational separation in entanglement theory](#), by Andi Gu, Salvatore F.E. Oliviero, and Lorenzo Leone, arXiv, April 2024 (43 pages).
- [1186] [Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels](#), by Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters, Physical Review Letters, March 1993 (5 pages).
- [1187] [Quantum Teleportation in a Nutshell](#), by Fabian Kössel, 2013 (35 slides).
- [1188] [Optimal ancilla-free Clifford+T approximation of z-rotations](#), by Neil J. Ross and Peter Selinger, arXiv, June 2016 (40 pages).
- [1189] [Quantum Teleportation is a Universal Computational Primitive](#), by Daniel Gottesman and Isaac L. Chuang, arXiv, August 1999 (6 pages).
- [1190] [Single-qubit gate teleportation provides a quantum advantage](#), by Libor Caha, Xavier Coiteux-Roy, and Robert Koenig, arXiv, November 2024 (37 pages).
- [1191] [Quantum gate teleportation between separated qubits in a trapped-ion processor](#), by Yong Wan, Daniel Kienzler, Stephen D. Erickson, Karl H. Mayer, Ting Rei Tan, Jenny J. Wu, Hilma M. Vasconcelos, Scott Glancy, Emanuel Knill, David J. Wineland, Andrew C. Wilson, and Dietrich Leibfried, arXiv, August 2019 (42 pages).
- [1192] [Applying an Evolutionary Algorithm to Minimize Teleportation Costs in Distributed Quantum Computing](#), by Leo Sünkel, Manik Dawar, and Thomas Gabor, arXiv, September 2024 (6 pages).
- [1193] [Deterministic teleportation of a quantum gate between two logical qubits](#), by Kevin S. Chou, Jacob Z. Blumoff, Christopher S. Wang, Philip C. Reinhold, Christopher J. Axline, Yvonne Y. Gao, L. Frunzio, M. H. Devoret, Liang Jiang, and R. J. Schoelkopf, Nature, September 2018  (preprint on [arXiv](#)).
- [1194] [Demonstrating advantages of dynamic quantum circuits on a hybrid superconducting qubit-cavity processor](#), by Hongbo Wu, Ling Hu, Jiasheng Mai, Munan Zhang, Libo Zhang, Yanyan Cai, Xiaowei Deng, Pan Zheng, Zhongchu Ni, Song Liu, Kun Fang, Dapeng Yu, and Yuan Xu, arXiv, August 2026 (17 pages).
- [1195] [Efficient Long-Range Entanglement Using Dynamic Circuits](#), by Elisa Bäumer et al., PRX Quantum, August 2023 (20 pages).
- [1196] [Methods for optimizing the synthesis of quantum circuits](#), by Timothée Goubault de Brugière, 2020 (213 pages).
- [1197] [Quantum Circuits for General Multiqubit Gates](#), by Mikko Mottonen, Juha J. Vartiainen, Ville Bergholm, and Martti M. Salomaa, arXiv, October 2004 (4 pages).
- [1198] [Exact two-qubit universal quantum circuit](#), by Jun Zhang, Jiri Vala, Shankar Sastry, and K. Birgitta Whaley, arXiv, May 2003 (5 pages).
- [1199] [Minimal Universal Two-qubit Quantum Circuits](#), by Vivek V. Shende, Igor L. Markov, and Stephen S. Bullock, arXiv, March 2004 (8 pages).
- [1200] [Optimizing resetting of superconducting qubits](#), by Ciro Micheletti Diniz et al., PRA, November 2023 (7 pages) .
- [1201] [Harnessing the Power of the Second Quantum Revolution](#), by Ivan H, November 2020 (13 pages).
- [1202] [Comparing the quantum switch and its simulations with energetically-constrained operations](#), by Marco Fellous-Asiani, Raphaël Mothe, Léa Bresque, Hippolyte Dourdent, Patrice A. Camati, Alastair A. Abbott, Alexia Auffèves, and Cyril Branciard, arXiv, February 2023 (21 pages).
- [1203] [Universal Quantum Computation via Superposed Orders of Single-Qubit Gates](#), by Kyrro Simonov, Marcello Calceffi, Jessica Illiano, Jacqueline Romero, and Angela Sara Cacciapuoti, arXiv, December 2024 (12 pages).
- [1204] [Experimental Aspects of Indefinite Causal Order in Quantum Mechanics](#), by Lee A. Rozema, Teodor Strömberg, Huan Cao, Yu Guo, Bi-Heng Liu, and Philip Walther, arXiv, July 2024 (22 pages).
- [1205] [Toward an Experimental Device-Independent Verification of Indefinite Causal Order](#), by Carla M.D. Richter, PRX Quantum, March 2026.
- [1206] [How many systems can be dephased before the quantum switch becomes causally definite?](#), by Yassine Benhaj, Kuntal Sengupta, and Cyril Branciard, arXiv, May 2026.
- [1207] [Improvement in quantum communication using quantum switch](#), by Arindam Mitra, Himanshu Badhani, and Sibasish Ghosh, arXiv, April 2023 (14 pages).
- [1208] [Quantum State Tomography via Compressed Sensing](#), by David Gross, Yi-Kai Liu, Steven T. Flammia, Stephen Becker, and Jens Eisert, Physical Review Letters, October 2010 (4 pages)  (preprint on [arXiv](#)).
- [1209] [How Many Shots Are Enough for a Quantum Circuit?](#), by Giuseppe Bisicchia, Alessandro Bocci, Ernesto Pimentel, and Antonio Brogi, arXiv, 2026 (57 pages).
- [1210] [StableShots: Online Shot Stopping for Quantum Circuit Execution](#), by Giuseppe Bisicchia, Alessandro Bocci, Ernesto Pimentel, and Antonio Brogi, arXiv, June 2026 (4 pages).
- [1211] [MIRAGE: Quantum Circuit Decomposition and Routing Collaborative Design using Mirror Gates](#), by Evan McKinney, Michael Hatridge, and Alex K. Jones, arXiv, August 2023 (13 pages).
- [1212] [Multi-qubit State Tomography with Few Pauli Measurements](#), by Xudan Chai, Teng Ma, Qihao Guo, Zhangqi Yin, Hao Wu, and Qing Zhao, arXiv, May 2023 (8 pages).
- [1213] [Certifying almost all quantum states with few single-qubit measurements](#), by Hsin-Yuan Huang, John Preskill, and Mehdi Soleimanifar, arXiv, April 2024 (63 pages).
- [1214] [Efficient tomography of microwave photonic cluster states](#), by Yoshiki Sunada, Yasunobu Nakamura, et al., arXiv, October 2024 (27 pages).
- [1215] [Scalable multiparticle entanglement of trapped ions](#), by H. Häffner, Rainer Blatt, et al., Nature, December 2005 (17 pages)  (preprint on [arXiv](#)).
- [1216] [Generation of multiphoton entangled quantum states by means of integrated frequency combs](#), by Christian Reimer, 2016 (7 pages)  (preprint on [repo.uni-hannover.de](#)).
- [1217] [Shadow Tomography of Quantum States](#), by Scott Aaronson, arXiv, November 2018 (29 pages).

- [1218] [Predicting many properties of a quantum system from very few measurements](#), by Hsin-Yuan Huang, Richard Kueng, and John Preskill, *Nature Physics*, June 2020 📄 (preprint on [arXiv](#)).
- [1219] [Triply efficient shadow tomography](#), by Robbie King, David Gosset, Robin Kothari, and Ryan Babbush, *arXiv*, April 2024 (42 pages).
- [1220] [Quantum process tomography via completely positive and trace-preserving projection](#), by George C. Knee, Eliot Bolduc, Jonathan Leach, and Erik M. Gauger, *arXiv*, January 2019 (13 pages).
- [1221] [Efficient Quantum Mixed-State Tomography with Unsupervised Tensor Network Machine Learning](#), by Wen-jun Li, Kai Xu, Heng Fan, Shi-ju Ran, and Gang Su, *arXiv*, August 2023 (7 pages).
- [1222] [An operational approach to quantum probability](#), by E. B. Davies and J. T. Lewis, *Communications in Mathematical Physics*, 1970 📄.
- [1223] [Quantum advantage in postselected metrology](#), by David R. M. Arvidsson-Shukur, Nicole Yunger Halpern, Hugo V. Lepage, Aleksander A. Lasek, Crispin H. W. Barnes, and Seth Lloyd, *Nature Communications*, July 2020.
- [1224] [Quantum and classical resources for unitary design of open-system evolutions](#), by Francesco Ticozzi and Lorenza Viola, *arXiv*, April 2017 (27 pages).
- [1225] [Quantum Nondemolition Measurements: The Route from Toys to Tools](#), by V. B. Braginsky and F. Ya. Khalili, *Reviews of Modern Physics*, January 1996 (11 pages) 📄.
- [1226] [Seeing a single photon without destroying it](#), by Gilles Nogues, Michel Brune, Jean-Michel Raymond, Serge Haroche, et al., *Nature*, July 1999 (4 pages) 📄.
- [1227] [Detecting an Itinerant Optical Photon Twice without Destroying It](#), by Emanuele Distante et al., *PRL*, June 2021 (6 pages).
- [1228] [Quantum non-demolition measurements in optics](#), by Philippe Grangier, Juan Ariel Levenson, and Jean-Philippe Poizat, *Nature*, December 1998 (7 pages) 📄.
- [1229] [Quantum non-demolition readout of an electron spin in silicon](#), by J. Yoneda, K. Takeda, A. Noiri, T. Nakajima, S. Li, J. Kamioka, T. Koder, and S. Tarucha, *Nature Communications*, March 2020.
- [1230] [Quantum Steering](#), by Roope Uola, Ana C. S. Costa, H. Chau Nguyen, and Otfried Gühne, *arXiv*, March 2020 (43 pages).
- [1231] [Quantum Steering on IBMQ](#), by Lennart Maximilian Seifert, Konstantin Beyer, Kimmo Luoma, and Walter T. Strunz, *arXiv*, January 2022 (11 pages).
- [1232] [Quantum Reservoir Engineering](#), by J.F. Poyatos, J.I. Cirac, and Peter Zoller, *Physical Review Letters*, February 1996 (14 pages) 📄.
- [1233] [Harnessing Disordered-Ensemble Quantum Dynamics for Machine Learning](#), by Keisuke Fujii and Kohei Nakajima, *Physical Review Applied*, August 2017 📄 (preprint on [arXiv](#)).
- [1234] [Quantum reservoir computing: a reservoir approach toward quantum machine learning on near-term quantum devices](#), by Keisuke Fujii and Kohei Nakajima, *arXiv*, November 2020 (13 pages).
- [1235] [Measurement, Dissipation, and Quantum Control with Superconducting Circuits](#), by Patrick Michael Harrington, 2020 (154 pages).
- [1236] [Reservoir Engineering using Quantum Optimal Control for Qubit Reset](#), by Daniel Basilewitsch, Francesco Cosco, Nicola Lo Gullo, Mikko Möttönen, Tapio Ala-Nissilä, Christiane P. Koch, and Sabrina Maniscalco, *arXiv*, September 2019 (11 pages).
- [1237] [Quantum reservoir engineering and single qubit cooling](#), by Mazyar Mirrahimi et al., 2013 (6 pages).
- [1238] [Stabilization of cat-state manifolds using nonlinear reservoir engineering](#), by Ivan Rojko, Matteo Simoni, Elias Zapusek, Florentin Reiter, and Jonathan Home, *arXiv*, July 2024 (31 pages).
- [1239] [Novel Technique for Robust Optimal Algorithmic Cooling](#), by Sadegh Raeisi, Mária Kieferová, and Michele Mosca, *arXiv*, June 2019 (5 pages).

This **Gate based quantum computing** section contains 152 bibliographical references and notes containing at least 5,965 pages.

5 Quantum computing engineering

After reviewing the basic principles of quantum physics and the mathematical basics of gate-based quantum computing, we will look here at the operational and physical operations of a quantum computer. We will describe the overall quantum computer architectures which depend closely on the characteristics of their qubits. We will look how errors are detected and corrected, and why and how quantum memories could be built. We'll assess the energy consumption of quantum computers and see whether it will be in the advantage or liability realm. We elaborate on the Lieb-Robinson, an underlooked description of how information and correlations propagate in entangled quantum systems, with several potential consequences on FTQC. We will also deal with the scientific and engineering uncertainties that may impact FTQC feasibility or, at least, their related roadmaps. Finally, we look at prospective economics of quantum computing from the cost and value perspective.

5.1 System architecture

5.1.1 Classical computing architecture

As a reminder and for future comparisons, here are the main components of a classical computer that you also find in various shapes and forms in smartphones, tablets, personal computers, game consoles and servers. Its key component is its microprocessor. It retrieves data and programs from a storage system and copies them to memory (RAM) entirely or on the fly as needed. The microprocessor then reads the program's instructions from memory in its cache to execute it one after the other and use conditional branching.

Data and programs can be retrieved remotely over a network and from remote servers on the Internet. The whole system is controlled by physical interfaces at input (keyboard, mouse, touchpad, joystick, webcam, microphones, scanners) and generates output (displays, audio, printers, other peripherals). The processor can be complemented by a graphics processor (GPU). It is either external to the microprocessor, for demanding requirements such as in computer aided design tools ala CATIA and Autodesk and video games or integrated into the microprocessor as is the case for almost all laptop and desktop processors. Depending on the configuration, the processor is surrounded by a variable number of external components that are soldered in the motherboard (Figure 230).

This is the case of the Intel chipsets like the Z390, which complements the core processors and manages a large part of the computer's inputs/outputs. Wi-Fi and cellular modems are associated with antennas. Of course, an internal and external power supply and a battery for mobile devices must be added. In the case of a server, you don't need all these end-user peripherals (mouse, keyboard, game pads, screens, speakers, etc).

On the power drain side, the processor and GPU heat up the most and require passive or active cooling depending on their power drain. In embedded systems such as smartphones, this is done with heat conducts and air. In PCs, it is supplemented by one or more fans.

In the most extreme cases, like with some servers, liquid cooling uses a water circuit to improve heat dissipation. One of the reasons why heat is generated by classical processing is the non-reversibility of classical computing. The Landauer bound is however a very distant floor: erasing one bit at 300 K costs $k_B T \ln 2 \approx 3 \times 10^{-21}$ J [25, 1240], some four to five orders of magnitude below the 10^{-16} J that a CMOS logic transition actually dissipates in charging and discharging capacitances. Irreversibility is a fundamental limit, not the reason current processors get hot.

One key difference between classical and quantum computers is the role of memory. In classical computing, processors access memory using some addressing scheme for retrieving both data and programs, they use a small number of registers and cache memory for computing. In quantum computing, gates act directly on qubit registers, which creates a resemblance with "in memory processing". But the qubit register is not classically addressable, its content cannot be read without measurement, and control and readout remain external to the register.

5.1.2 Gate-based quantum computers key parameters

The key performance indicators of gate-based quantum computers were created by **David DiVincenzo** when he was an IBM researcher (Figure 231). He first listed them in 1996 [1241] and formalized them in his 2000 paper [1242].

While individual qubits barely existed, he listed a set of requirements for a universal gate-based quantum computer, described below. These are a useful checklist for the gate model and not a necessary and sufficient condition across modalities. Measurement-based, analog and continuous-variable architectures satisfy several of them only in a modality-dependent way, and error-corrected architectures reinterpret others, starting with initialization and measurement.



Figure 230: Illustrative personal computer hardware architecture, shown as a familiar baseline rather than as a complete model of contemporary classical computing. It leaves aside accelerators, the full memory hierarchy, high-speed interconnects and the distributed system features that matter when a quantum computer is integrated in a data center. © Olivier Ezratty, 2021-2024.




Figure 231: Selected gate-model requirements, known as the DiVincenzo criteria. They describe the gate-based circuit model and are neither necessary nor sufficient across all modalities, with documented exceptions for measurement-based, analog and continuous-variable architectures. © Olivier Ezratty, 2021-2024, inspired by a slide from Pascale Senellart.

Well-characterized qubits. Quantum computers use qubits that exploit quantum objects that can have two distinct and measurable states. Their physical characteristics are well known. The architecture is scalable in the sense that it can exploit many physical qubits and then, logical qubits relying on these physical qubits and quantum error correction codes.

Initializable qubits. In general, to the value $|0\rangle$ often called “ground state” for the associated quantum objects, corresponding, for example, to the lowest energy

level of an elementary particle or an artificial atom as for superconducting qubits.

Coherence times. It must be greater than quantum gates activation times. The time during which the qubits are in a coherent state must be greater than the quantum gates activation time in order to be able to execute an algorithm containing a sufficiently long sequence of quantum gates. Error correction codes using a large number of physical qubits have the benefit of extending this usable computing time.

Universal quantum gates set. The quantum hardware must allow the creation of a universal gate set. It depends on the qubit technology. It requires a minimum set of single-qubit gates allowing the creation of any rotation in the Bloch sphere, completed by a CNOT two-qubits gate.

Measurement. With the ability to measure qubits state at the end of computing, which seems obvious. This measurement should not influence the state of other qubits in the system. Ideally, the measurement error rate should be well below 0.1%, a bar no platform meets today: median readout error sits around 1 to 2% on superconducting devices and 0.1% on the best trapped ion and spin qubit setups.

David DiVincenzo added two other optional criteria that are used instead for quantum communications:

Flying qubits conversion. The ability to map the state of a stationary qubit onto a propagating photonic or electronic mode and back. Transduction, quantum state transfer and physical transport are three distinct operations that are often merged.

Transport these moving qubits. from one point to another reliably and remotely. This will allow to manage quantum telecommunications, distributed architectures of quantum computers and to set up *blind computing* architectures allowing to distribute treatments while protecting their confidentiality. The technology will quickly become essential to enable the distribution of quantum computations over several quantum processors, a bit like we do with multi-core chips or with processing distribution architectures over several CPUs and several servers. Some vendors like IonQ have announced that they will rely on this architecture. This will be useful for qubit architectures that will be limited in the number of qubits, which may only be able to consolidate a few hundred at most. It will thus be necessary to be able to link remote processor qubits and keep them entangled. Different quantum interconnection techniques are possible. The most generic is optical and is not much constrained by distance. At rather short distances, microwave links are possible, particularly to couple superconducting qubits, as well as shuttling electrons.

DiVincenzo's criteria are quite basic. From a practical and operational point of view [1243], quantum computers can also be characterized by another set of parameters as follows:

Qubit number. The relevant metric is the number of qubits per quantum processor. It conditions the available computing power in memory space, which enables potential speedups. As this power theoretically increases exponentially with the number of qubits, it is a key parameter. As of September 2026, the largest superconducting quantum processors available in the cloud was IBM's 156-qubit Heron r2, announced in November 2024, while the largest commercially accessible neutral atom array reaches 256 qubits with QuEra. IBM's 1,121-qubit Condor was never made publicly accessible due

to faulty qubits and the system being a testbed for scalable control electronics. Exposed qubit count says nothing about fidelity, connectivity or usable circuit volume, which must be quoted separately with their measurement protocol and date.

Some technologies are easier to miniaturize and scale than others. It is necessary to integrate in this miniaturization both the quantum qubit chips and the elements that control them. On top of that, we must ensure that decoherence and noise does not increase as the number of qubits is growing. After the qubit number limit per processor is reached, vendors need to find ways to quantumly interconnect multiple processors, using various means, mainly microwave guides (for superconducting qubits) and otherwise photonic links. These interconnect solutions create entanglement resources between multiple quantum processors.

Qubit fidelities. When executing quantum gates and reading their state, qubit fidelity conditions the ability to execute long algorithms. It has a direct impact on the supported algorithm depth. It also impacts the capacity to run quantum error correction codes and create logical qubits with an arbitrary fidelity level. Fidelities are characterized for qubit initialization and reset, single and two qubit gates as well as for qubit readout, the most demanding one being two-qubit gate fidelities. Trapped ions qubits have the best fidelities of all modalities, at 99.92% for all-pairs two-qubit gates on Quantinuum's 98-qubit Helios in 2025 and 99.97% and, in 2025 as well, above 99.99% with a two qubit gate on a two-qubit test chip from Oxford Ionics [1244, 1245]. Superconducting qubits scale-up better, with median two-qubit gate errors now sitting at 3×10^{-3} on IBM Heron-class devices and below 10^{-3} on the best two-qubit gates. Google Willow also crossing the surface code threshold with similar fidelities [1246]. Fidelity uniformity on a large chip remains the harder problem. Cold atom qubits scale better and reach 99.5% to 99.7% for two-qubit gates [1247]. Electron spins qubits could scale best in theory and have caught up faster than expected, exceeding 99% for one- and two-qubit gates on devices fabricated in a 300 mm foundry [1248], with academic devices above 99.5%.

Qubit operations (QOPS) describe the capability of a quantum computer in number of operations, i.e., a number of quantum gates. This metric is usually in powers of 10 and is now frequently used in vendors and even country roadmaps, like in the UK. One benefit of this metric is it can work both in NISQ and FTQC regimes and smoothen the transition between these. It can lead to use the KISQ, MISQ, GISQ and TISQ scales corresponding to a number of instructions in 10^3 , 10^6 (*aka megaquop*), 10^9 (*aka gigaquop*) and 10^{12} (*aka teraquop*) and ISQ meaning "instruction scale quantum".

Qubits connectivity. It conditions the quantum algorithms execution speed. The greater this physical connectivity, the faster the code execution will be. With a low connectivity, the quantum code compiler will have




Figure 232: Quantum computers datacenters integration requires some care in different domains. On-premise D-Wave annealers currently fit in a $3\text{ m} \times 3\text{ m} \times 3\text{ m}$ packaging. Most other NISQ quantum computers occupy a similar space, the larger being IBM System Two which is about 6.7×6.7 meters wide and 4.6 meters high. According to vendor plans, FTQC quantum computers will have variable sizes, from the existing size of NISQ systems to huge plants even not fitting in a data center like with PsiQuantum. In the EuroHPC data centers hosting NISQ systems, quantum computers are usually installed in separate rooms supporting a low-noise environment (about 50 dB in the room [1249]) and stable temperature at $\pm 1^\circ\text{C}$ between 18°C and 25°C for superconducting qubits [1249]. © Olivier Ezratty, 2021-2026. See [1243] for more information on QPU integration in HPC center.

to add a lot more operations to link the qubits together, particularly relying on SWAP gates, that also add their own noise in the process. This connectivity varies greatly from one technology to another. In 2D technologies, as with superconducting and silicon qubits, it is currently limited to neighboring qubits (1 to 2 and 1 to 3 in the case of IBM until 2025 with Eagle/Heron processors, and 1 to 4 with Google processors since Sycamore in 2019 and IBM processors since Nighthawk that was announced in 2025). IBM Nighthawk contains 120 qubits and 218 tunable couplers on a 12×10 square lattice, with a 1-to-4 nearest neighbor connectivity. The long-range “c-couplers” that make non-local qLDPC codes practical [1250] show up on a separate Loon test chip. There is also a proposal to generalize a 1 to 3 triangle connectivity [1251]). It seems that connectivity is better with some types of trapped ion qubits although with some limitations depending on the implementation and the range, and with cold atoms which can be moved around in 2D space for implementing two-qubit gates with two arbitrary atoms (Figure 233) [1252].

One proposal to overcome nearest neighbor connectivity overhead was made in 2025 by Wolfgang Dür from the University of Innsbruck [1253]. It uses data qubits and entanglement qubits in 2D layouts, where entanglement can be generated between all pairs of entanglement qubits using long-range Bell states. The cost is that all entanglement qubits have to be initially entangled as a large cluster state which may be hard to implement at scale. The proposal contains another approach that doesn’t require the preparation of such cluster states, but with a higher time overhead. Everything has a cost!

Qubit parallel operation. How qubit gates can be parallelized over different qubit zones without disruption

will also condition the speed of execution of quantum algorithms.

Execution time. It depends on multiple factors like gates and measurement times, error correction overhead as well as qubit connectivity and fidelities. Fast quantum computers with fast but lower fidelity gates are not necessarily faster than those with slow qubits with better fidelities and qubit connectivity [1254]. Some vendors, particularly in the photonic qubits realm, are mentioning the notion of clock rate, which in their case, corresponds to the frequency of single photon emissions. With most matter qubits (superconducting, ions, atoms), gate operations have variable durations. Most of the time, single-qubit gates are faster than two-qubit gates. Also, in FTQC regime, which will be covered later in this part of the book, one critical constraint comes from qubit measurement and error syndrome decoding time.

Operating temperature. For the processor and their equipment which is very dependent on the type of qubits. The Holy Grail is of course to operate at room temperature. The currently operational quantum computers based on superconductors operate at a very low temperature of 15 mK (1 mK = 1 millikelvin, 0 kelvin = -273.15°C), but some types of qubits still in the research stage are supposed to operate at room temperature, such as those based on photons and NV centers (point defects in diamond pairing a substitutional nitrogen atom with an adjacent lattice vacancy, like with Quantum Brilliance and SaxonQ). However, this is not necessarily the case for associated equipment such as photon sources and detectors for photon qubits. Operating at very low temperature is a way to preserve the coherence of the qubits. But the lower the temperature, the smaller the energy that can be radiated by the qubits and their con-




Figure 233: Qubit connectivity depends on the device generation and not only on the modality. The graph shown for each vendor is the native interaction graph of its September 2026 flagship, and reconfigurable platforms such as neutral atoms change that graph during a circuit run. You need to minimize the overhead of SWAP gates when creating CNOT gates between arbitrary qubits. The caveat with ions and atoms is that their many-to-many connectivity comes with a price: gates are very slow compared to circuit-based qubits. Things have changed in the superconducting qubits realm. IBM used for a while a 1-to-2 and 1-to-3 limited connectivity with its heavy-hex architecture, to limit the effects of crosstalk while Google used a nearest-neighbor 1-to-4 connectivity. In 2025, IBM switched gear and with the Nighthawk chip implemented a 1-to-4 nearest neighbor connectivity on a square lattice. Non-local qLDPC error correction will rely on the long-range c-couplers demonstrated on the separate Loon test chip. © Olivier Ezratty, March 2025.

trol electronics. Operating qubits at 100 mK or 1K, like with electron spin qubits, creates a much larger available cooling budget to control the qubits than operation at 15 mK. Neutral atoms and trapped ions are frequently said to operate at room temperature. As a matter of fact, they are always cooled at very low temperature, mostly using laser beams and their enclosing chamber is also frequently cooled at 4K. The system itself must be installed in stable temperature environment with a usual tolerance of $\approx 1^\circ\text{C}$.

Total energy consumption. We will investigate this and study it in a global manner with incorporating all quantum computer components: the processor itself, all its control electronics and photonics as well as the involved cryogenic systems. A power figure only means something with a good definition of its perimeter. Using the wall-plug power of the complete system, including cryogenics, control electronics and colocated classical computers is recommended. On that basis, mixing vendor statements and site measurements, quantum computers had in mid-2025 a power drain sitting between 2 kW and over 200 kW depending on the qubit type and number of qubits. Comparing two systems requires aligning utilization, the share of classical compute inside perimeter and the energy per completed task, not only the instantaneous power. And this is going to be somewhat different with FTQC QPUs.

System rackability. How will quantum computers be deployed in data centers? Does it fit in standard rack systems? It is notably planned by the startup Pasqal, as

well as for Quandela's photon generators and LightOn's optical processors, as well as micro-wave external electronics from companies like Zurich Instruments and Qblox. Alpine Quantum Technologies from Austria also announced in 2021 was fitting its trapped ion computing in two standard 19-inch racks [1255]. It is associated with issues of weight, space, cooling, and power supply. What kind of fluids must be used for cooling, usually cold water, connected to the first stage compressor of cryostats, whatever their size? Quantum computers must also withstand the usual data centers conditions like vibrations, dust, and electromagnetic environment, or be separated in special isolated facilities (Figure 232). Superconducting qubit systems have specific electromagnetic environment constraints (Table 6). They could sit in the modular building blocks used in the most recent data centers. These last three operational parameters play a role when deploying computers or quantum accelerators in data centers. It plays a critical role since, for most applications, quantum computers will be offered through cloud services.

All these considerations to gauge the capabilities of a quantum computer involve the discipline of quantum computers benchmarking! As Kristel Michiels points out, benchmarks can be used when the number of qubits is below 50 when comparing the rendering of algorithms between quantum computers and their emulation on supercomputers [1256]. Beyond that, it will be more difficult.

Benchmarked quantum computers will generally have dissimilar characteristics: different universal quantum gates requiring compilers to assemble different quantum gates to execute the same algorithm, and different error correction codes, adapted to the qubit fidelities, circuit size and the primary quantum gate set of the compared computers. The dissimilarities will be much greater than between two Intel and AMD processors or two smartphone chips!

Parameter	Requirement
DC magnetic field	< 100 μ T for each of the axes.
AC magnetic field	< 1 μ T peak-to-peak spectrum amplitude for each of the axes at frequency range 5 Hz – 1000 Hz.
Floor vibrations	< 400 μ m/s RMS spectrum amplitude for each of the axes at frequency range 1 Hz – 200 Hz, or equivalently, ISO vibration limit for office spaces.
Sound pressure	< 80 dBA when integrated over the frequency range 20 Hz – 20 kHz.
Temperature	$\Delta T < \pm 1$ °C within 12 hours around any set point between 20 – 25 °C.
Humidity	25 – 60%, non-condensing

Table 6: Physical constraints for the installation of an IQM quantum computer at the Angers Bull/Eviden datacenter. Source: slide 7 in [1257].

5.1.3 Quantum computers segmentation

There is not just one category of quantum computers, but many. We must at least distinguish gate-based quantum computers and analog computers, including quantum annealing computers such as the ones from D-Wave.

But there are at least six categories of quantum computing paradigms as shown in Figure 234:

Quantum emulators are used to execute quantum algorithms on traditional computers ranging from simple laptops to supercomputers, depending on the number of qubits and algorithm depth and precision to be emulated. It is based on large vectors and matrices computing. Code emulation is used to test quantum algorithms without quantum computers. Quantum emulators are sometimes called quantum simulators, but this name should be avoided to prevent confusion with... analogue quantum simulators. These are quantum processors simulating quantum physics phenomena, for example magnetism or the tridimensional structure of molecules. Quantum emulators may however also simulate qubit noise model like Atos/Eviden QLM emulator. We can make a distinction between an exact digital simulation and approximate digital simulation, emulating a digital error rate that is equal or below NISQ hardware. This can help simulate a greater number of qubits. Quantum emulators can also reproduce the (quantum) physical characteristics of various qubits and in that case, they

also implement some form of digital quantum simulation. To date, supercomputers can fully emulate up to the equivalent of 40 to 50 qubits [1258]. Records have however been broken with more than 100 qubits, with a low number of quantum gates and using various techniques like tensor compression. Emulating quantum computers requires a lot of power both on the memory side, to store 2^N quantum register states for N qubits, and for the associated processing that relies on floating-point matrix multiplications.

Quantum inspired computing is about using classical algorithms running on classical hardware that are inspired by quantum algorithms and bring some new efficiencies. They are not about emulating quantum code on a classical computer. Typical quantum inspired algorithms use tensor network libraries and techniques like MPS and DMRG.

Quantum annealers draw on the adiabatic theorem, which uses a slow and controlled evolution of a set of qubits linked together according to a particular topology (“Pegasus” or “Zephyr” in the case of D-Wave). The process is first initialized in the ground state of the Hamiltonian and the adiabatic theorem guarantees the convergence of the system towards a low energy state, ideally the ground state. This technique is used to search for an energy minimum to solve various problems (simulations, optimizations, machine learning). The coefficients of the Hamiltonian are the couplings (weights of the interactions between qubits) and the self-couplings (weights of the qubits) and the variables of an instance are the spins of each qubit. Many problems can be translated into quantum annealing problems using QUBO or Ising problem formulations. Strictly speaking the theorem only guarantees convergence for a closed system evolving slowly enough compared to the inverse squared minimum gap. Real annealers are open systems coupled to a thermal bath at a temperature well above their gap for most of the anneal, so what they perform is better described as thermally assisted quantum annealing than as adiabatic quantum computation [1259]. D-Wave seems to bring interesting gains in computation time for some use cases, but this is strongly disputed by many experts, particularly in the case of combinatorial optimizations.

Quantum simulators or analogue quantum computers work in an analog and not digital way, with continuous parameters linking the qubits together. The most common technique is based on neutral atoms, like with Pasqal and QuEra. It is also labelled “Rydberg quantum annealing” [1260]. Trapped ions [1261–1263], superconducting qubits [1264], spin qubits [1265], photon qubits [1266, 1267] and other various quantum systems can also be used for running quantum simulations [1266], but no commercial vendor seems to promote this paradigm when they can also implement gate-based quantum computing which is supposed to be more generic [1268]. Some vendors like ParityQC, IQM and Kipu Quantum are still proposing architectures mixing digital and ana-




Figure 234: The different computing paradigms with gate-based quantum computers, analog quantum computers (grey), and classical systems for code emulation or quantum inspired tensor network-based solutions (in blue). This paradigm segmentation is a September 2026 view and is not exhaustive. Some vendors like Pasqal and QuEra explore more than one paradigm on the same hardware, such as digital and analog modes on neutral atoms. © Olivier Ezratty, 2022-September 2026.

log quantum computing platforms. Quantum annealers and quantum simulators are sometimes labelled “analog NISQ” systems since they are as noisy as gate-based NISQ systems, also named “digital NISQ”.

Digital quantum computers *aka* universal quantum computers use qubits with quantum gates capable of executing all sorts of quantum algorithms. They are also labelled general purpose quantum computers. It is a concept associated with universal gate-based quantum computing, in opposition to analog computing which is said to have a more limited set of applications. It is also linked to an economic rationale. Indeed, computing markets usually develop naturally with generic platforms, not too narrowly specific ones. However, the “digital” may be a misnomer here for these types of computers. Although algorithms are decomposed in series of quantum gates, these are still analog computers per se. Their internal quantum state is analog. Their mathematical transformations are analog. The quantum systems are affected by analog noise. It is a key reason why they are complicated to create.

Gate-based superconducting quantum computers are currently limited to 156 qubits (IBM Heron). Cold atoms are at 1,600 with Inflektion’s Sqale and 1,180 with Atom Computing since 2023. These sizes correspond to the ability to position individual atoms in space with optical tweezers, but not for the implementation of gate operations.

Qubit quantum noise is detrimental to computing and requires the usage of logical qubits made of many physical qubits and quantum error correction codes (QEC). While waiting for these fault-tolerant quantum computers to ramp up with logical qubits, we are using

non corrected qubits in the so-called NISQ for “Noisy Intermediate-Scale Quantum”, an expression coined by John Preskill in 2017-2018 [1269]. It describes existing and future general-purpose quantum computers supporting 50 to a few hundred physical qubits. These can run algorithms with a limited circuit depth due to the qubit error rates like variational quantum circuits. Their performance is improved by using quantum error suppression and quantum error mitigation techniques. They are supposed at some point to exceed supercomputers computing capacities for solving specific problems. Then, much later, we will have fault tolerant (FTQC) quantum computers, with a very large number of physical qubits and over 100 logical qubits, which will really open the realm of useful quantum computing.

There are now several other variations of universal quantum computers that deserve some description:

Continuous variables quantum computers, or analog quantum computers with universal gates, replace the two-level qubit by a *qumode*, a bosonic mode whose information is carried by the continuous and unbounded quadratures $\hat{q}$ and $\hat{p}$ in an infinite-dimensional Hilbert space [1271, 1272]. Their information content is sometimes counted in *qunats*, the natural-logarithm analog of the qubit. This category of quantum computing was proposed in 1999 by Seth Lloyd and Samuel L. Braunstein [1273]. They are usually based on continuous variable photons but other qubit types like trapped ions are used (Figure 235).

MBQC, or Measurement Based Quantum Computing, is an architecture adapted to flying qubits and particularly to photon qubits which can’t easily be entangled to create two qubit gates. The process creates a set of

Direct Variable vs Continuous Variable encoding of quantum information




Figure 235: Direct variable and continuous variable encoding of quantum information. Inspired from [1270]. 2020-2025.

entangled qubits at the beginning of computing, aka cluster states [1274]. It is followed by qubits readouts in an ordered way, enabling the implementation of traditional gates. MBQC also implements some massive parallelism, adapted to the limited and finite processing depth of flying qubits. PsiQuantum plans to use a variant of this technique named FBQC [1275]. Quandel has also plans to implement MBQC using its own proprietary cluster states generation technique.

Q Analog or analogue quantum simulation?

You can debate whether using the “analog” or “analogue” term to describe these systems as described in [1276]. The authors state that “*to fully appreciate the distinction between analog and digital computation, it is crucial to distinguish the physical processes underlying a computation on the level of the hardware and the computational model itself*”.

Analog means continuous input and output encoding together with continuous control parameters whereas “analogue” means “equivalent”, a physical system (the simulator) being equivalent to another physical system (the one we want to simulate). So an “analogue quantum simulator” can be an “analog quantum simulator”.

To make things simple, we will keep using “analog” in this book.

Topological quantum computing is based on specific anyon qubits that are self-corrected [1277]. The low-level programming model of these qubits is much different from universal quantum computers. This is the path chosen by Microsoft. Its development seems to be quite sluggish. This technology still requires error correction codes and fault-tolerance mechanisms although with a

lower overhead. Google and researchers in China are studying the use of superconducting qubits to reproduce the mechanism of braiding and non-abelian systems to implement topological qubits [1278].

Another segmentation of these models can be made with two dimensions: discrete or continuous data encoding and discrete or continuous variables computing given the vendor position is rough, some being positioned in various slots (Pasqal also wants to do gate-based computing) [1279] (Figure 236).

Quantum Accelerator. It is a quantum computer used as a complement to a supercomputer or HPC, usually to run variational algorithms like VQE (Variational Quantum Eigensolvers) combining a classical part that prepares the data structure that feeds a quantum accelerator [1280]. The quantum computer serves as an accelerator for the HPC which can be a node or the whole HPC, using CPUs and/or GPUs/TPUs. GPUs/TPUs are themselves also accelerators for the CPUs. There are some design issues requiring tight integration between the HPC and the quantum computer, particularly with regards to batch loading and to the way the quantum algorithm is executed multiple times (Figure 237). A quantum computer contains itself at least a classical computer. It converts digital signals (gates) into analog signals (the microwaves or lasers controlling the qubits and handling their readout). This quantum computer classical computer will need to be as close as possible to the HPC computing capacities to improve the turnaround, depending on data I/O throughput and latency requirements. It may lead to create custom designs integrating an HPC and one or several quantum accelerators [1281, 1282]. Other quantum accelerator designs contain more or less generic higher-level software layers with connectors driving various quantum and classical architectures (annealers, gate-based, emulators) [1283]. Even in the FTQC regime, many use cases will require the use of HPCs to prepare the quantum circuits used by quantum computers. It will be for example the case for chemical simulations which require




Figure 236: Discrete vs continuous data encoding vs data processing. Inspired from [1279]. Updated in September 2026.




Figure 237: basics of a hybrid classical/quantum computing hardware architecture, given that, in all cases, quantum computing is always hybrid. In this scheme, a quantum computer consumes a quantum circuit and generates some useful results that are post-processed either in the quantum computer or in a outside computer which can be an HPC if the workload mandates it. At the lowest level, the quantum computer contains a classical computer that compiles the circuit, controls the electronics/photonic to implement gates and readout, handles real-time quantum error correction and fault tolerance (FTQC) or quantum error mitigation (NISQ), and error detection + suppression (both in FTQC and NISQ). An external system like an HPC can be used to prepare the quantum circuit with some classical preprocessing and potentially run some results post-processing coming from the quantum computer. With NISQ variational algorithms, a classical optimizer refines the variational quantum circuit, which, for performance reasons, is better off running in a classical computer colocated within the quantum computer. In most cases, the quantum computer is considered to be an specialized accelerator of a classical computer. © Olivier Ezratty, 2021-September 2026.

a classically prepared approximation, in the case of a quantum phase estimation algorithm.

This inventory is only an appetizer. We will have the opportunity to detail these various architectures.

5.1.4 Exotic quantum computing paradigms

We are always in for many surprises and new programming paradigms proposed by academic researchers, that nearly nobody in the ecosystem seems to be evaluating or even understanding. Here is an inventory of some

of these models with their country of origin, the year when it was proposed, and an evidence status.

A good segmentation should use five levels: theoretical proposal, component experiment, few-qubit demonstration, programmable prototype and commercial access. But nearly all the models below sit at the first level. They are therefore segmented here with separating first, the variations and extensions of existing gate-based models, and second, the really stretched models which require new physics.

- Holonomic quantum computation** (Italy, 1999, evidence status: few-qubit demonstrations in NMR, trapped ions and superconducting circuits) is a scheme for implementing quantum gates using non-Abelian geometric phases, called holonomies, instead of the usual dynamical evolution driven by the energy of a Hamiltonian. This paradigm elaborated by Paolo Zanardi and Mario Rasetti [1284] encodes logical information in a degenerate eigenspace of a parameter-dependent Hamiltonian, and then drives the control parameters slowly around a closed loop in parameter space. It is in principle insensitive to some types of errors. Related work was published in 2000 by Jonathan A. Jones, Vlatko Vedral, Artur Ekert and Giuseppe Castagnoli, with an NMR implementation [1285], followed by a trapped ion proposal from Luming Duan, Ignacio Cirac and Peter Zoller [1286] and with superconducting circuits [1287]. It then pivoted in 2012 with Erik Sjöqvist from Sweden to non-adiabatic holonomic quantum computation (NHQC) [1288]. It led to dark path holonomic qudit computation in 2022 [1289, 1290] which uses non-adiabatic geometric phases to implement robust and efficient logic gates on qudits, while avoiding unwanted state transitions through specially designed evolution paths known as “dark paths”. The model enhances noise resilience and scalability, enables faster gate operations and more efficient use of quantum resources compared to traditional qubit-based schemes. Another variant is the Brachistochrone Non-Adiabatic Holonomic Quantum Control (B-NHQC) created in China in 2020 [1291] with a demonstration implemented with trapped ions in 2026 [1292]. In quantum systems, a brachistochrone is a time-optimal evolution from one quantum state to another [1293].
- Meromorphic quantum computing** (UK, 2026, evidence status: theoretical proposal) was proposed by Simon Burton and Hussain Anwar from Quantinuum [1294]. It uses a mathematical framework that translates the standard rules of quantum mechanics from linear algebra into a language of geometry. It represents quantum states as points on a sphere. Quantum circuits including errors correction codes, are described as “meromorphic functions”, which are algebraic fractions. Computing stable stationary points of these functions helps handle error correction efficiently.
- Quantum frequential computing** (France, 2024, evidence status: theoretical proposal) [1295] which optimizes quantum control mechanisms, ensuring a minimum quadratic runtime speedup.
- Pauli Quantum Computing** (China, 2024, evidence status: theoretical proposal) [1296], where instead of using vector states like $|0\rangle$ and $|1\rangle$, quantum state are expanded as sums of Pauli operators such as I , X , Y , and Z . Operators like I and X symbolically take on the role of the usual basis states, replacing $|0\rangle$ and $|1\rangle$. It relies on density matrices to interpret these Pauli elements. Algorithms like amplitude estimation and quantum search (*ala* Grover) could be executed more efficiently than in traditional gate-based frameworks where they bring only a quadratic speedup. It still requires specific control strategies and optimized measurement protocols.
- Enhancing quantum processor capabilities during idle times** (Austria, 2026, evidence status: theoretical proposal), a model proposed by Wolfgang Dür [1297]. It proposes to create entangled states (d dimensional cluster state or dD-CS) with ancilla qubits during computing idle times or before the start of a circuit run. These are then consumed by the data qubit register using gate teleportation. It is an intermediate between measurement based quantum computing (MBQC) and gate-based computing.
- Autonomous Quantum Machine (AQM)** (USA, 2025, evidence status: theoretical proposal) is a proposed technique to sequence a series of gates in specific hardware and control drive for cold atoms, trapped ions and superconducting qubits [1298]. The gate sequence is physically hardcoded into some spatial or temporal layout in the system during a classically controlled initialization phase before the AQM is launched. With trapped ions, the ions roll down a potential gradient defined by different stationary continuous laser beams. As ions pass through a single beam, this runs a single-qubit gate, while two ions passing through a beam simultaneously execute some two-qubit gate. The circuit is defined by the physical arrangement of these beamed zones along the ion slide. With Rydberg atoms, the gate sequence is timed and transmitted using Passively Mode-Locked Lasers (PMLs) acting as autonomous quantum clocks. They emit stable, ultrashort pulses. Other lasers driving some sequence of gates are emitted through a delay line. The mechanism with superconducting qubits is using some quantum heat engine that extracts work from a thermal gradient between a hot and cold baths and drives some gate. This mechanism is a theoretical construct that doesn't scale yet to real-life circuits.
- Higher-order quantum computing (HOQC)** (Japan, 2025, evidence status: theoretical proposal with small photonic component experiments) is a

computing paradigm where computing is about transforming some quantum operations (expressed as unitaries) into other quantum operations [1299, 1300]. Practical use cases could be the implementation of arbitrary unitary inversion, transposition, and storage-and-retrieval (SAR), a kind of “quantum circuit” memory.

- **Fuzzy-processing quantum computation** (China, 2026, evidence status: theoretical proposal) proposed by Yan-Xiong Du [1301], which is a quantum equivalent to the classical “fuzzy logic” paradigm created in 1965 by Lotfi Zadeh from UC Berkeley (1921-2017) [1302, 1303].

And now, the stretch physical models which require new hardware engineering, which are most of the time, borderline science-fiction.

- **Molecular-based holographic quantum computing** (Denmark, 2008, evidence status: theoretical proposal) [1304] is explored by Molecular Quantum Solutions (Denmark). It uses a single mesoscopic ensemble of trapped polar molecules, in which hundreds of qubits are encoded as collective excitations carrying distinct spatial phase patterns. Each pattern is addressed by optical Raman processes, while gates and readout transfer the qubit state to a stripline microwave cavity coupled to a single Cooper pair box, where the actual two-level dynamics happens.
- **Time-crystal quantum computing** (Poland and Australia, 2025, evidence status: theoretical proposal) [1305] uses time crystals quantum systems that spontaneously break time-translation symmetry to enhance coherence and stability in quantum processors. Time crystals can be implemented with various proposed physical substrates including cold atoms and micro-mirrors. It enables more reliable and noise-resilient quantum operations. On paper. It could help build quantum neural networks and variational algorithms.
- **Lorentz quantum computing** using hyperbolic bits (China, 2021, evidence status: theoretical proposal) [1306, 1307]. Qubits are associated with “hybits” whose quantum evolution follows complex Lorentz transformations rather than the standard unitary evolution of conventional quantum mechanics. It supports a universal set of logical gates. A MIS (maximum independent set) search algorithm designed for this model could be potentially run in polynomial time.
- **Non-Hermitian quantum computer** (NQC) (China, 2025, evidence status: theoretical proposal) which uses a non-unitary quantum gate G and proposes to solve NP-complete problems in polynomial times. It was created by the same authors as Lorentz quantum computing [1308].
- **Cooling-based Quantum Computation** (CQC) (South Korea, 2024, evidence status: theoretical proposal) [1309] looks like replacing adiabatic quantum computation with active cooling of the system to drive it to its minimum energy level, and solve problems more efficiently, particularly optimization problems.
- **Symplectic computing** (Russia, 2024, evidence status: theoretical proposal) [1310] replaces the unitary evolution of qubits with symplectic transformations, using the quantum-symplectic duality between complex and real formulations of the Schrödinger equation. It is using real-valued symplectic transformations instead of complex unitary ones. No hardware implementation is proposed. It is not clear whether this model really uses quantum phenomena like entanglement.
- **Metamaterials based quantum computing** (USA-Colombia, 2026, evidence status: theoretical proposal) is an architecture proposal that combines symmetry protection of relative-motion qubits (ions, atoms, superconducting, and semiconductor based) via the generalized Kohn theorem, control via twisted light carrying orbital angular momentum (OAM), which couples to the motional and correlation degrees of freedom and is distinct from the photon spin angular momentum used in ordinary polarization-driven transitions, and metamaterial nanofocusing using Weyl semimetal plasmonics [1311].
- **Quantum computing using quantum gravity superposition** (Israel, 2024, evidence status: theoretical proposal), with space-time warping enabling computing in superposed orders and bringing an exponential-exponential speedup [1312, 1313]. The actual physical implementation is not described in these papers.
- **Unruh-DeWitt qubits** (UDW) (Austria and Canada, evidence status: theoretical proposal) [1314], two-level detectors coupled to a relativistic quantum field, whose gates are mediated by the field itself. The link to black holes is indirect and goes through the Unruh effect, where an accelerated observer sees a thermal bath behind an acceleration horizon analogous to a black hole horizon.
- **Quantum gravity topological** phenomenon-based quantum computing (USA and UAE, evidence status: theoretical proposal) [1315].

Most of these paradigms are purely theoretical, with no practical implementation proposal, and, hopefully, no related startup.

5.1.5 Qubit types

Quantum computers using qubits are devices that encode information into one physical property or observable of some quantum object that can have two possible

mutually exclusive states, that can be initialized, modified with quantum gates, and then measured.

They are sometimes individual quantum objects, as with atoms (trapped ions and cold atoms), electrons (quantum dots silicon qubits) or photons! And only one at a time! In the case of superconducting qubits or Majorana fermions, the quantum state is based on a large number of electrons arranged in Cooper pairs that share the same quantum state, the pairs of electrons that are created at superconducting temperature. With NV centers and some exotic qubits, qubits are constructed with ensembles of quantum objects or with heterogeneous quantum objects like mixing electron spins and atom nuclear spins.

Qubits can also be classified in two meta-breeds: stationary or moving (flying) as in Figure 238. Those based on trapped ions, cold atoms, electrons spin, NV centers and superconducting loops are stationary. Flying qubits are based on photons that physically circulate from quantum gate to quantum gate as well as on flying electrons. They move around from a source, through physical devices implementing quantum gates and land on detectors. In all cases, the quantum gates are dynamically activated by electronic circuits or lasers and operate on the qubits where they are (stationary qubits) or in the path of their transit (flying qubits).

I now refer you to the large quantum computing hardware section of this book to learn more about the various types of qubits that are coming out of research laboratories and industry vendors.

The various figures of merit of these qubit architectures are:

Qubits stability which is evaluated with measured relaxation and dephasing parameters: the energy relaxation time T_1 , the Ramsey dephasing time T_2^* which includes inhomogeneous broadening, and the echo-corrected T_2 . These are decay constants of specific experiments and not the lifetime of superposition or of entanglement. Associated with quantum gate times, gate and idle error rates, leakage and crosstalk, they condition the number of quantum gates that can be used in a quantum circuit. Trapped ions have the longest coherence times to date, in the seconds range and beyond on small registers, with the caveat that these values are measured on few qubits and currently degrade as the register grows.

Qubits fidelity is related to the errors level that is evaluated with single and two qubit gates as well as with initialization and readout. Any evaluation must account for the used metrics, the register size, the benchmark protocol, its date, and on whether the number is a best pair, a median, or a value measured under simultaneous operation. On two-qubit gate fidelity measured on small registers, trapped ions lead as of September 2026, while superconducting and neutral atom systems report their best numbers on much larger registers.

Qubits connectivity is the way they are linked together, which conditions many aspects such as quantum circuit execution speed, its supported depth, and even the types of error correction codes that can be used. Trapped ions and neutral atoms provide all-to-all or reconfigurable connectivity, which lowers the SWAP overhead but comes with gate times three to four orders of magnitude longer than superconducting qubits, and with scaling limitations as of 2026. Connectivity can change over time and technology generations within a given qubit modality.

Large scale entanglement, without being limited to the immediately neighboring qubits. So far, nobody does it well.

Operating temperature and for the accompanying electronics. NV centers can operate their qubits at ambient temperature at a limited scale, while superconducting qubits require about 15 mK. In general, all qubit types require some form of cryogenics, usually ranging from 10 mK to 4K.

Qubits density and their control electronics which impacts scalability. This rather favors quantum dots electron spin qubits.

Manufacturing process which depends on many parameters. In the case of neutral atoms, for example, it is not necessary to create specialized circuits to control the qubits, whereas it is necessary for most other qubit technologies.

Scalability potential which depends on many systems parameters, both at the fundamental level with the qubit stability and fidelities at large scale but also with the various enabling technologies. Unfortunately for your forecasting, scalability potentials do not align with qubits technologies present maturities!

The level of qubits is evolving rapidly. It is described in this excellent document from the German cybersecurity agency (Figure 239) [1316]. It mentions other technologies not listed in this inventory.

Figure 240 shows another way to put it [1147, 1317]. It segments the types of qubits according to three dimensions: the clock frequency of the quantum gates (roughly, the gates number that can be executed per second), the number of operations before errors occur, and the quantum gates fidelity (separating the one- and two-qubit gates). These last two axis are roughly homothetic because the number of operations before errors are generated depends on the error rate.

Trapped ions have better gates than superconducting qubits but are quite slow. Silicon qubits are for the moment quite fast, at least, as fast as superconducting qubits. Neutral atoms are slower in gate-based mode. A last axis is missing: the number of qubits as of today and technology scalability. The chart reflects 2018 data presented in 2019 and is now outdated for most qubit types, with superconducting and spin qubits hav-




Figure 238: Separating stationary and flying qubits. © Olivier Ezratty, 2021-2025.




Figure 239: Degree of maturity of various qubit technologies from the scalability standpoint, (L) flags modalities for which at least one encoded logical qubit has been demonstrated. The label does not say which milestone was reached, and these are distinct criteria: encoded state, repeated correction cycles, break even against a stated physical baseline, below-threshold scaling with distance, logical gates, and fault-tolerant operation. Scaling related to existing platforms but not on potential to scale beyond the existing state of the art. Inspired from the 2020 version of [1316]. © Olivier Ezratty, July 2024-August 2025.

ing gained roughly an order of magnitude in operations before error since then.

5.1.6 Architecture overview

We will provide here an overview of the general architecture of a quantum computer, using the example of a superconducting qubit accelerator (Figure 241 and Figure 242).

First, a bit like some external GPUs, quantum computers are implemented as co-processors or accelerators of

classical computers that power and control them. A quantum computer is always driven by a classical computer, as can be a GPU for video games or for training neural networks in deep learning. These conventional computers are used to run the programs that drive the quantum processor with physical operations to be performed on the qubits and are interpreting qubits readout results.

The classical computer closely controls the operation of the quantum computer by triggering at a precise rate the operations on the qubits that are performed




Figure 240: Comparison of qubit computing depth and gate speed, in 2018. Source: [1317].

by various electronic devices creating various electronics and photonic signals controlling quantum gates and quantum readout. It takes into account quantum gates execution time and the known qubits coherence time, i.e., the time during which the qubits remain in a state of superposition and entanglement.

In addition to its classical control computer, our quantum computer includes at least the components labeled from 1 to 6 that we will describe one by one, first with an overview below, then later, with a more detailed view. The other types of quantum computers have similarities and differences that we will mention whenever relevant.

Quantum registers are collections of qubits. IBM's Osprey QPU, with its 433 superconducting qubits of which 413 were exposed in the `ibmq_seattle` processor when it went live in May 2023, held the largest exposed superconducting register at that date. It was retired in September 2023. Quantum registers hold the quantum state that the QPU coherently evolves, exploiting superposition and entanglement.

Quantum gate controllers are physical devices that act on the quantum register qubits, both to initialize them and to perform quantum gates on them. These gates are applied iteratively, according to the quantum circuit to be executed. They can also be used to manage error correction codes. Data and instructions are not separated as they are with classical microprocessors, since the gate sequence is applied directly to the qubits.

Readout of qubit states is used to obtain the result at the end of the sequential execution of a quantum circuit, and to evaluate error syndromes during quantum error correction, including in mid-circuit. This cycle of initialization, calculation and measurement is usually applied

several times, each repetition being a shot or run. The shots estimate the quantity the quantum circuit was designed to produce, which can be either a bitstring distribution or an expectation value. The signals coming from qubit readout are converted into digital values and transmitted to the conventional computer which controls the system and implements post-processing tasks to interpret results. In common cases, such as with D-Wave and IBM, computing is repeated at least a couple thousand times. The reading devices are connected to their control electronics via superconducting wires in the case of superconducting computer qubits.

Quantum chip usually includes quantum registers, quantum gates controls and measuring devices. These are fed by various electronic or photonic signals coming from outside the chip. Devices can be highly heterogeneous for the types of qubits using lasers for initialization, quantum gates and qubit measurement like with trapped ions and cold atoms. Current chips are not very large. Their larger current size is in centimeters. Each qubit is relatively large, a transmon footprint being a few hundred microns on a side for superconducting qubits, or down to about 100 nm for electron spin qubits, whereas leading CMOS logic nodes are labelled 3 nm or 2 nm, a commercial name that no longer maps onto any physical dimension, with actual gate lengths closer to 15 to 20 nm. Superconducting chips are usually integrated in an OFHC (Oxygen-Free High thermal Conductivity) copper packaging, purified and freed from oxygen precisely to maximize its thermal conductivity at low temperature, so that the chip is properly thermalized with the mixing chamber plate. This package is fitted with coaxial connectors so that it can be fed by the microwaves controlling qubit gates. The I/O count




Figure 241: Typical high-level architecture of a gate-based quantum computer. Most qubits are hosted in a cryogenic environment for multiple reasons: to ensure their operation and avoid ambient thermal noise, to improve vacuum quality (for trapped ions and cold atoms) and to ensure the good operations of single photon sources and detectors (for photon qubits). The qubit register hosts the quantum data while quantum gates act on this quantum memory and quantum readout measure the qubit value. All this is driven by a mix of cryogenic and ambient temperature control electronics and photonics depending on the modality. By design, a quantum computer contains some classical computers which can control the rest of the system. A minimal quantum computer is a system which can run a quantum circuit that is provided as an input and yield some post-processed results as an output. In the case of FTQC systems, this includes all the error correction classical tasks which must be run in real-time and therefore are part of the quantum computer. An external computing system like an HPC can be dedicated to the generation of the quantum circuit to run on the system and for some heavy classical post-processing tasks. © Olivier Ezratty, 2020-2026.

scales with the qubit count and quickly becomes the binding constraint: about 160 coaxial lines for the 53-qubit generation of 2019, and well over a thousand for a 156-qubit Heron r2, which is why flex cabling and cryo-multiplexing are now on every vendor roadmap. These line counts depend on the multiplexing scheme and on how much of the control electronics sits inside the cryostat.

The chip package is integrated into two small concentric magnetically insulated enclosures, in aluminum and in Cryoperm, a high-nickel cryogenic alloy trademarked by **Vacuumschmelze** and fabricated into shields by companies like **MuShield** and **Amuneal**.

Cryogeny usually keeps the qubit chip and its surrounding control electronics at a temperature close to absolute zero. It contains part of the control electronics and the quantum chip(s) to avoid generating disturbances that prevent the qubits from working. The Holy Grail would be to operate qubits at room temperature but the corresponding architectures such as in NV centers are not yet operational and there are still practical performance reasons to operate it at low-temperatures. A dry dilution refrigerator cools its 50 K and 4 K stages with a pulse tube fed by compressed helium 4, then circulates a closed helium 3 and helium 4 mixture that is condensed, separated in the still, and diluted in the mixing chamber, the dilution of helium 3 into the helium

4 rich phase providing the cooling power below about 100 mK. The pulse-tube and mixture compressors are themselves cooled by a facility water loop, similar to the compressors used in classical air conditioning. Other types of qubits use cryostats in different places: with cooling photon sources or detectors in photon qubits systems, or for cooling ultra-vacuum pumps with cold atoms.

Control electronics in the cryostat enclosure and at room temperature. The qubit control electronics drive the physical devices used to initialize, modify, and read the qubit status. With superconducting qubits, quantum gates are activated with microwave pulses using frequencies between 4 and 8 GHz and generally located outside the cryostat. These microwaves are transmitted on coaxial electrical wires between their source and the quantum processor, with superconducting cables below 4K. Their generators still take up a lot of space. They are not very miniaturized at this stage. Interesting work aims at integrating these microwave generators and readers inside the cryostat enclosure, if only to limit the wiring. These are frequently based on cryo-CMOS technology, CMOS components that are tailored to work at low temperature, 4K for many and as low as 20 mK for some. Figure 242 provides a rough representation of an entire superconducting qubits based quantum computer.




Figure 242: Typical physical components of a superconducting qubit quantum computer. It contains a classical computer that drives the whole system. A chip is cooled at about 15 mK at the bottom of a dry dilution refrigerator. The 50 K and 4 K stages are cooled by a pulse tube fed with compressed helium 4 in a closed circuit. Below 4 K, cooling comes from a closed circulating helium 3 and helium 4 mixture, condensed on the still and separated into a helium 3 rich phase and a dilute phase, the dilution across the phase boundary in the mixing chamber providing the cooling power at base temperature. The pulse-tube and mixture compressors are themselves cooled by a facility water loop with a secondary water-air exchanger. Qubits are driven by a mix of microwave pulses and other electrical signals coming through various cabling from room temperature electronics. These electronic systems are themselves driven by a classical computer. © Olivier Ezratty, 2020-2026.

Q Digital or analog?

Is a gate-based quantum computer a digital or an analog computing device? I'd argue it is always an analog computing system that is wrapped in a digital shell.

Quantum gates from the Clifford group ("digital group") are insufficient to generate an exponential speedup, since a Clifford-only circuit on stabilizer inputs is classically simulable in polynomial time by the Gottesman-Knill theorem [1318]. To obtain it, we need more freedom to twist our register vector state in any position from its Hilbert space ("analog"). This is obtained by creating arbitrary rotation single qubit gates ("analog") in NISQ regime. In a way, a variational circuit with these rotation gates is "analog". Qubit control enables the implementation of these "analog" gates.

However, when switching to a fault-tolerant system, we are becoming more "digital". Arbitrary rotation gates can't be applied directly to encoded

qubits. Clifford gates are the easy case, implemented transversally or, in surface codes, by lattice surgery between code patches [1319]. The non-Clifford gates, T and Toffoli, are the expensive ones and go through the strange and cumbersome mechanism of magic state distillation and injection [1320]. Arbitrary rotations ("analog") are therefore first approximated by long sequences of Clifford and T gates ("digital"), based on the Solovay-Kitaev theorem, and in practice with the Ross-Selinger synthesis which reaches a near-optimal T-count of about $3 \log_2(1/\epsilon)$ [1321].

In a way, a quantum algorithm is analog with its arbitrary rotations used for example in a QFT or inverse QFT, but is approximated in a "digital" mode due to error correction. And in all cases, the readout being classical bit, from the outside, a quantum computer is a digital system. But again, since results are random, you sample it a large number of times and make a sort of average of the results, which turns you back in some analog of continuous mode.

Classical computers are part of a quantum computer to handle many tasks that are usually split over several of these. One computer connects the whole system to the outside world, with, for example, some cloud or Intranet infrastructure. One may manage quantum com-

puting jobs with batch management tools. One contains various middleware and software tools and at least a code compiler, transpiler and optimizer. Another one controls the QPU control electronics, themselves driving the qubits, whatever the qubit modality. [1322] makes

a theoretical proposal for an autonomous Quantum Processing Unit (aQPU) that does not require a classical computer. One motivation would be to contain energy consumption, but the model practicality is not obvious to say the least. In this model, a quantum computer contains a quantum program and a quantum clock to time the quantum gates.

Classical computers run quantum error mitigation and/or quantum error correction tools, unless these are directly handled by specific hardware in the control electronics side of the QPU. These classical computing resources can be an HPC, like a Nvidia DGX server appliance, or several regular servers depending on the quantum processor capacity. Finally, a supercomputer or HPC may be needed to handle quantum algorithms pre- and post-processing tasks. They may be quite heavy in the case of quantum chemistry. In that case, the quantum computer is a sort of accelerator or quantum coprocessor for the HPC or supercomputer.

QEC/FTQC sizing. As we have seen in the part dedicated to FTQC in 5.2.11, error correction and fault-tolerance adds its resource burden that creates new challenges for system architecture. They are manifold and deal with sizing the number of supported physical qubits and their classical control, the related high-capacity cryogenics, potential internal quantum processor interconnect devices, and all the classical qubit readout and error correction chain [1327]. The significant error correction classical workload justifies to place error correction syndrome detection tasks support within the quantum computer, in order to reduce its latency. In some systems like with silicon spin qubits, this part of the quantum computer may consume more energy than the rest of the system [1328].

5.1.7 Processor layout

To better understand the previous explanation, here is a chip layout with 8 superconducting qubits, from ETH Zurich shown in Figure 243. Although it is already a few years old, the underlying concepts are generic.

- **Qubits** are in the white rectangles. These are tiny Josephson effect superconducting circuit loop.
- **Coupling circuits** link them together. It is used to control entanglement between pairs of qubits.
- **Single-qubit gates** use the blue and purple contacts. It sends microwaves to the qubits. These pins are powered via cables by very high frequency current sources, sending microwaves photons, between 4 and 8 GHz. These frequencies must be different between adjacent qubits of the same circuit to avoid crosstalk. It is the combination of these frequencies that will trigger different types of quantum gates and entanglements between adjacent qubits.

- **Measurement** takes place with other circuits, also fixed in the component. In superconducting qubits, readout is dispersive: each qubit is coupled to a readout resonator, and a microwave probe tone reflected off that resonator comes back with a phase and amplitude that depend on the qubit state. A Purcell filter protects the qubit from decaying into the readout line, and the outgoing signal is amplified by a quantum-limited parametric amplifier before leaving the cryostat through superconducting coaxial lines [472].

Qubits must interact with each other but as little as possible with their environment until measurement. This is one of the reasons why they are usually cooled to a temperature close to absolute zero and magnetically isolated from the outside. The choice of materials for the chips also plays a role in minimizing the noise that could affect the qubits and bring them out of their coherent state.

The circuit shown in Figure 244 shows the relationship between qubit gates time and coherence time during which the qubits remain stable. The number of two-qubit gate cycles is only partly limited by the qubit coherence time, gate, idle, measurement, leakage and correlated errors setting the executable depth together with it. In the current generations of superconducting qubit quantum computers, two-qubit gates last about 90 ns, against T_1 and T_2 of roughly 170 and 130 μ s on a Heron r2. Coherence alone would allow well over a thousand gate cycles. But median two-qubit error rates around 3×10^{-3} cap the useful depth at a few hundred sequential entangling gates, so the circuit is damaged long before coherence runs out. This is why two-qubit error rates matter more than coherence times (T_1 and T_2). Coherence times are much longer for quantum computers using trapped ions, in the seconds range, but their gate times are three to four orders of magnitude longer as well. In silicon quantum dots qubits, gate times are short, tens of nanoseconds, and coherence is long only once dynamical decoupling is applied: the raw T_2^* sits around 40 μ s while the echoed T_2 reaches a few milliseconds. Algorithms must therefore optimize the number of gate cycles to be executed, which is furthermore constrained by the physical connectivity between qubits.

In diagrams describing quantum algorithms, such as the one in Figure 244, the double bar after measuring the state of a qubit indicates that a normal bit has been recovered, at 0 or 1. The physical register is reused from the beginning to the end of a run, but the input and output widths in qubits of a quantum circuit do not necessarily match. There may be fewer used qubits in the output readout due to the use of various ancilla qubits or lost qubits.




Figure 243: Small 8-qubit superconducting processor from ETH Zurich created in 2017 showing its various lines controlling the qubits. A single line is used for qubit readout thanks to microwave frequency multiplexing. You can observe that the resonators in blue and green for each qubit have different numbers of loops. The couplers in orange are used to drive two qubit gates. Sources: [1323] and the thesis [1324].





Figure 244: Timeline showing the relation between useful computing time, gate times, gate coherence times and qubit operation errors taking the example of superconducting qubits. In most cases, qubit error rates have a significant toll on the circuit execution and fidelity long before we reach the T_1 and T_2 of a quantum processor. One reason is the way these coherence times are evaluated. In the current NISQ regime, the number of usable gate cycles depends on the device, the circuit structure, the concurrency of operations and error detection, suppression and mitigation techniques. As of mid-2026, a superconducting circuit was run on an IBM system with 452 two-qubit-gate layers using 62 physical qubits by Q-CTRL [1325] and a trapped-ion circuit with 1,225 gate layers using 50 qubits on Quantinuum system by Jülich SC [1326]. Deeper circuits didn't generate results distinguishable from random noise. © Olivier Ezratty, 2020-September 2026.



5.1.8 Bibliography and notes

These are the bibliographical references and notes for the **System architecture** subsection.

- [25] [Experimental verification of Landauer’s principle linking information and thermodynamics](#), by Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz, *Nature*, March 2012 (3 pages)  (preprint on [arXiv](#)).
- [472] [A Quantum Engineer’s Guide to Superconducting Qubits](#), by Philip Krantz, William D. Oliver, et al., *Applied Physics Review*, June 2019 (67 pages).
- [1147] [Introduction to Quantum Computing](#), by William Oliver, MIT, December 2019 (21 slides).
- [1240] [Irreversibility and Heat Generation in the Computing Process](#), by Rolf Landauer, *IBM Journal of Research and Development*, July 1961 (9 pages) .
- [1241] [Topics in Quantum Computers](#), by David P. DiVincenzo, *arXiv*, December 1996 (20 pages).
- [1242] [The Physical Implementation of Quantum Computation](#), by David DiVincenzo, *arXiv*, 2000 (9 pages).
- [1243] [First Practical Experiences Integrating Quantum Computers with HPC Resources: A Case Study With a 20-qubit Superconducting Quantum Computer](#), by Eric Mansfield, Laura Schulz, et al., *arXiv*, September 2025 (9 pages).
- [1244] [Scalable, High-Fidelity All-Electronic Control of Trapped-Ion Qubits](#), by Clemens M. Löschnauer, Jacopo Mosca Toba, Alex C. Hughes, Sebastian A. King, Marius A. Weber, Raghavendra Srinivas, René Matt, Rustam Nourshargh, David T. C. Allcock, Christopher J. Ballance, Chris Matthiesen, Maciej Malinowski, and Thomas P. Harty, *PRX Quantum*, October 2025.
- [1245] [Trapped-ion two-qubit gates with >99.99% fidelity without ground-state cooling](#), by Alex C. Hughes, Raghavendra Srinivas, Clemens M. Löschnauer, Hannah M. Knaack, René Matt, Christopher J. Ballance, Maciej Malinowski, Thomas P. Harty, and Robin T. Sutherland, *arXiv*, October 2025.
- [1246] [Quantum error correction below the surface code threshold](#), by Rajeev Acharya, Michael Newman, Austin G. Fowler, Hartmut Neven, and Julian Kelly, *Nature*, February 2025.
- [1247] [High-fidelity parallel entangling gates on a neutral-atom quantum computer](#), by Simon J. Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H. Li, Alexandra A. Geim, Tout T. Wang, Nishad Maskara, Harry Levine, Giulia Semeghini, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin, *Nature*, October 2023.
- [1248] [Industry-compatible silicon spin-qubit unit cells exceeding 99% fidelity](#), by Paul Steinacker, Nard Dumoulin Stuyck, Wee Han Lim, Tuomo Tantt, Andre Saraiva, Chih Hwan Yang, and Andrew S. Dzurak, *Nature*, October 2025.
- [1249] [Bringing the Cat home - Installing on-prem superconducting cat qubits](#), by Rémi de la Vieuville, TQCI International Feedback on First QPU integration, December 2025 (25 slides).
- [1250] [High-threshold and low-overhead fault-tolerant quantum memory](#), by Sergey Bravyi, Andrew W. Cross, Jay M. Gambetta, Dmitri Maslov, Patrick Rall, and Theodore J. Yoder, *Nature*, March 2024.
- [1251] [A New Layout-Aware Quantum Circuit Synthesis with Davio Lattices Using Triangular Connectivity](#), by Sonia Yang, Ali Al-Bayat, and Marek Perkowski, *arXiv*, June 2025 (13 pages).
- [1252] [A quantum processor based on coherent transport of entangled atom arrays](#), by Dolev Bluvstein, Harry Levine, Giulia Semeghini, Tout T. Wang, Sepehr Ebadi, Marcin Kalinowski, Alexander Keesling, Nishad Maskara, Hannes Pichler, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin, *Nature*, April 2022.
- [1253] [Long-ranged gates in quantum computation architectures with limited connectivity](#), by Wolfgang Dür, *arXiv*, July 2025 (11 pages).
- [1254] [Assessing Finite Scalability in Early Fault-Tolerant Quantum Computing for Homogeneous Catalysts](#), by Yanbing Zhou, Shangjie Guo, et al., *arXiv*, November 2025 (20 pages).
- [1255] [Compact Ion-Trap Quantum Computing Demonstrator](#), by Ivan Pogorelov, Rainer Blatt, Philipp Schindler, Thomas Monz, et al., *PRX Quantum*, June 2021 (23 pages).
- [1256] [Benchmarking gate-based quantum computers](#), by Kristel Michielsen, Madita Nocon, Dennis Willsch, Fengping Jin, Thomas Lippert, and Hans De Raedt, *arXiv*, June 2017 (33 pages).
- [1257] [Feedback on IQM Spark 5 Integration at Eviden’s HPC Lab](#), by Francois Archereau and Hermanni Heimonen, TQCI International Feedback on First QPU integration, December 2025 (14 slides).
- [1258] [Massively parallel quantum computer simulator, eleven years later](#), by Hans De Raedt, Fengping Jin, Dennis Willsch, Madita Willsch, Naoki Yoshioka, Nobuyasu Ito, Shengjun Yuan, and Kristel Michielsen, *Computer Physics Communications*, April 2019.
- [1259] [Adiabatic quantum computation](#), by Tameem Albash and Daniel A. Lidar, *Reviews of Modern Physics*, January 2018 (64 pages)  (preprint on [arXiv](#)).
- [1260] [Solving optimization problems with local light shift encoding on Rydberg quantum annealers](#), by Kapil Goswami, Rick Mukherjee, Herwig Ott, and Peter Schmelcher, *arXiv*, December 2023 (18 pages).
- [1261] [Non-equilibrium critical scaling and universality in a quantum simulator](#), by A. De, P. Cook, K. Collins, W. Morong, D. Paz, P. Titum, G. Pagano, A. V. Gorshkov, M. Maghrebi, and C. Monroe, *arXiv*, September 2023 (10 pages).
- [1262] [Interaction graph engineering in trapped-ion quantum simulators with global drives](#), by Antonis Kyprianidis, A. J. Rasmusson, and Philip Richerme, *arXiv*, October 2023 (27 pages).
- [1263] [Realization of programmable Ising models in a trapped-ion quantum simulator](#), by Yao Lu, Wentao Chen, Shuaining Zhang, Kuan Zhang, Jialiang Zhang, Jing-Ning Zhang, and Kihwan Kim, *arXiv*, November 2023 (11 pages).
- [1264] [Observation of Inelastic Meson Scattering in a Floquet System using a Digital Quantum Simulator](#), by Ziting Wang, Franco Nori, Heng Fan, et al., *arXiv*, August 2025 (3 pages).
- [1265] [Analog Quantum Simulation of the Dynamics of Open Quantum Systems with Quantum Dots and Microelectronic Circuits](#), by Chang Woo Kim, John M. Nichol, Andrew N. Jordan, and Ignacio Franco, *arXiv*, October 2022 (20 pages).
- [1266] [Realizing tight-binding Hamiltonians using site-controlled coupled cavity arrays](#), by Abhi Saxena, Arnab Manna, Rahul Trivedi, and Arka Majumdar, *Nature Communications*, August 2023  (preprint on [dx.doi.org](#)).
- [1267] [A hybrid-frequency on-chip programmable synthetic-dimension simulator with arbitrary couplings](#), by Xiao-Dong Zeng, Guang-Can Guo, et al., *arXiv*, August 2025 (11 pages).
- [1268] [A scalable superconducting quantum simulator with long-range connectivity based on a photonic bandgap metamaterial](#), by Xueyue Zhang, Eunjong Kim, Daniel K. Mark, Soonwon Choi, and Oskar Painter, *arXiv*, August 2022 (34 pages).
- [1269] [Quantum Computing in the NISQ era and beyond](#), by John Preskill, *arXiv*, July 2018 (20 pages).
- [1270] [Sub-Universal Models of Quantum Computation in Continuous Variables](#), by Giulia Ferrini, June 2018 (35 slides).
- [1271] [Universal Quantum Computing with Arbitrary Continuous-Variable Encoding](#), by Hoi-Kwan Lau and Martin B. Plenio, *arXiv*, May 2016 (5 pages).

- [1272] [Continuous-variable quantum computing in the quantum optical frequency comb](#), by Olivier Pfister, 2019 (16 pages)  (preprint on [arXiv](#)).
- [1273] [Quantum computation over continuous variables](#), by Seth Lloyd and Samuel L. Braunstein, arXiv, October 1998 (9 pages).
- [1274] [A One-Way Quantum Computer](#), by Robert Raussendorf and Hans J. Briegel, Physical Review Letters, May 2001 (4 pages) .
- [1275] [Fusion-based quantum computation](#), by Sara Bartolucci, Patrick Birchall, Hector Bombín, Hugo Cable, Chris Dawson, Mercedes Gimeno-Segovia, Eric Johnston, Konrad Kieling, Naomi Nickerson, Mihir Pant, Fernando Pastawski, Terry Rudolph, and Chris Sparrow, Nature Communications, February 2023.
- [1276] [Analogue Quantum Simulation: A New Instrument for Scientific Understanding](#), by Dominik Hangleiter, Jacques Carolan, and Karim P. Y. Thébault, arXiv, March 2023 (187 pages).
- [1277] [Non-Abelian Anyons and Topological Quantum Computation](#), by Chetan Nayak, Steven H. Simon, Ady Stern, Michael Freedman, and Sankar Das Sarma and, arXiv, March 2008 (73 pages).
- [1278] [Non-Abelian braiding of Fibonacci anyons with a superconducting processor](#), by Shibo Xu, Dong-Ling Deng, et al., arXiv, March 2024 (37 pages).
- [1279] [Quantum computing using continuous-time evolution](#), by Viv Kendon, arXiv, April 2020 (19 pages).
- [1280] [Quantum Accelerators for High-Performance Computing Systems](#), by Keith A. Britt, Fahd A. Mohiyaddin, and Travis S. Humble, arXiv, December 2017 (7 pages).
- [1281] [Quantum Accelerator Stack: A Research Roadmap](#), by K. Bertels, A. Sarkar, A. Krol, R. Budhrani, J. Samadi, E. Geoffroy, J. Matos, R. Abreu, G. Gielen, and I. Ashraf, arXiv, May 2021 (39 pages).
- [1282] [QPU-System Co-Design for Quantum HPC Accelerators](#), by Karen Wintersperger, Hila Safi, and Wolfgang Maurer, arXiv, September 2022 (15 pages).
- [1283] [Quantum Computer Architecture: Towards Full-Stack Quantum Accelerators](#), by K. Bertels, A. Sarkar, A.A. Mouedenne, T. Hubregtsen, A. Yadav, A. Krol, and I. Ashraf, arXiv, September 2019 (20 pages).
- [1284] [Holonomic quantum computation](#), by Paolo Zanardi and Mario Rasetti, Physics Letters A, December 1999 (5 pages)  (preprint on [arXiv](#)).
- [1285] [Geometric quantum computation using nuclear magnetic resonance](#), by Jonathan A. Jones, Vlatko Vedral, Artur Ekert, and Giuseppe Castagnoli, Nature, February 2000  (preprint on [arXiv](#)).
- [1286] [Geometric Manipulation of Trapped Ions for Quantum Computation](#), by L.-M. Duan, J. I. Cirac, and P. Zoller, Science, June 2001  (preprint on [arXiv](#)).
- [1287] [Non-Abelian Holonomies, Charge Pumping, and Quantum Computation with Josephson Junctions](#), by Lara Faoro, Physical Review Letters, January 2003  (preprint on [arXiv](#)).
- [1288] [Non-adiabatic holonomic quantum computation](#), by Erik Sjöqvist, D M Tong, L Mauritz Andersson, Björn Hessmo, Markus Johansson, and Kuldip Singh, New Journal of Physics, October 2012 (11 pages).
- [1289] [Dark path holonomic qudit computation](#), by Tomas André and Erik Sjöqvist, Physical Review A, December 2022 (7 pages).
- [1290] [Parallel transport in rotating frames and projective holonomic quantum computation](#), by Ole Sönnerborn, arXiv, May 2025 (7 pages).
- [1291] [Brachistochrone Non-Adiabatic Holonomic Quantum Control](#), by Bao-Jie Liu, Zheng-Yuan Xue, and Man-Hong Yung, arXiv, February 2021 (5 pages).
- [1292] [Experimental Demonstration of a Brachistochrone Nonadiabatic Holonomic Quantum-Gate Scheme in a Trapped Ion](#), by Xi Wang, Hui Ren, L.-N. Sun, K.-F. Cui, J.-T. Bu, S.-L. Su, L.-L. Yan, and G. Chen, arXiv, March 2026 (6 pages).
- [1293] [Time-Optimal Quantum Evolution](#), by Alberto Carlini, Physical Review Letters, February 2006  (preprint on [arXiv](#)).
- [1294] [Meromorphic Quantum Computing](#), by Simon Burton and Hussain Anwar, arXiv, May 2026 (17 pages).
- [1295] [Quantum Frequential Computing: a quadratic runtime advantage for all computations](#), by Mischa P. Woods, arXiv, March 2025 (102 pages).
- [1296] [Pauli quantum computing: \$I\$ as \$|0\rangle\$ and \$X\$ as \$|1\rangle\$](#) , by Zhong-Xia Shang, arXiv, December 2024 (3 pages).
- [1297] [Enhancing quantum processor capabilities during idle times](#), by Wolfgang Dür, arXiv, June 2026 (14 pages).
- [1298] [Proposals for experimentally realizing \(mostly\) quantum-autonomous gates](#), by José Antonio Marín Guzmán, Yuxin Wang, Tom Manovitz, Paul Erker, Norbert M. Linke, Simone Gasparinetti, and Nicole Yunger Halpern, arXiv, November 2025 (5 pages).
- [1299] [Higher-order quantum computing with known input states](#), by Vanessa Brzić, Satoshi Yoshida, Mio Murao, and Marco Túlio Quintino, arXiv, November 2025 (88 pages).
- [1300] [Higher-Order Quantum Operations](#), by Philip Taranto, Simon Milz, Mio Murao, Marco Túlio Quintino, and Kavan Modi, arXiv, March 2025 (65 pages).
- [1301] [Fuzzy-processing quantum computation](#), by Yan-Xiong Du, arXiv, June 2026 (6 pages).
- [1302] [Fuzzy sets](#), by Lotfi A. Zadeh, Information and Control, June 1965 (16 pages).
- [1303] [Fuzzy Logic and Soft Computing—Dedicated to the Centenary of the Birth of Lotfi A. Zadeh \(1921–2017\)](#), by Sorin Nădăban, Mathematics, September 2022 (3 pages) .
- [1304] [Holographic quantum computing](#), by Karl Tordrup, Antonio Negretti, and Klaus Molmer, arXiv, July 2008 (4 pages).
- [1305] [Time-tronics: from temporal printed circuit board to quantum computer](#), by Krzysztof Giergiel, Peter Hannaford, and Krzysztof Sacha, arXiv, May 2025 (15 pages).
- [1306] [Lorentz Quantum Computer](#), by Wenhao He, Zhenduo Wang, and Biao Wu, arXiv, June 2021 (10 pages).
- [1307] [The Power of Lorentz Quantum Computer](#), by Qi Zhang and Biao Wu, arXiv, September 2024 (18 pages).
- [1308] [Physics and Computation: A Perspective From Non-Hermitian Quantum Computer](#), by Qi Zhang and Biao Wu, arXiv, June 2025 (9 pages).
- [1309] [Quantum Computation by Cooling](#), by Jaeyoon Cho, arXiv, June 2024 (8 pages).
- [1310] [Quantum vs. Symplectic Computers](#), by Igor Volovich, arXiv, July 2024 (9 pages).
- [1311] [Symmetry-Protected Quantum Computing using Metamaterials](#), by Neil F. Johnson, Ferney J. Rodriguez, and Luis Quiroga, arXiv, May 2026 (9 pages).
- [1312] [Quantum Algorithms in a Superposition of Spacetimes](#), by Omri Shmueli, arXiv, May 2024 (59 pages).
- [1313] [Unitary Closed Timelike Curves can Solve all of NP](#), by Omri Shmueli, arXiv, October 2024 (23 pages).
- [1314] [A Universal Quantum Computer From Relativistic Motion](#), by Philip A. LeMaitre, T. Rick Perche, Marius Krumm, and Hans J. Briegel, arXiv, May 2025 (5 pages).
- [1315] [Topological QBits in Flux-Quantized Super-Gravity](#), by Hisham Sati and Urs Schreiber, arXiv, November 2024 (16 pages).
- [1316] [Status of quantum computer development](#), by BSI, BSI, January 2025 (222 pages).

- [1317] [When could NISQ algorithms start to create value in discrete manufacturing?](#), by Oxana Shaya, arXiv, September 2022 (63 pages).
- [1318] [The Heisenberg Representation of Quantum Computers](#), by Daniel Gottesman, arXiv, July 1998 (20 pages).
- [1319] [Surface code quantum computing by lattice surgery](#), by Dominic Horsman, Austin G. Fowler, Simon Devitt, and Rodney Van Meter, New Journal of Physics, December 2012 (28 pages).
- [1320] [Universal quantum computation with ideal Clifford gates and noisy ancillas](#), by Sergey Bravyi and Alexei Kitaev, Physical Review A, February 2005 (14 pages)  (preprint on [arXiv](#)).
- [1321] [Optimal ancilla-free Clifford+T approximation of z-rotations](#), by Neil J. Ross and Peter Selinger, Quantum Information and Computation, September 2016 (53 pages) .
- [1322] [Autonomous Quantum Processing Unit: What does it take to construct a self-contained model for quantum computation?](#), by Florian Meier, Marcus Huber, Paul Erker, and Jake Xuereb, arXiv, January 2024 (18 pages).
- [1323] [The European Quantum Technologies Roadmap](#), by Antonio Acín, Immanuel Bloch, Harry Buhrman, Tommaso Calarco, Christopher Eichler, Jens Eisert, Daniel Esteve, Nicolas Gisin, Steffen J. Glaser, Fedor Jelezko, Stefan Kuhr, Maciej Lewenstein, Max F. Riedel, Piet O. Schmidt, Rob Thew, Andreas Wallraff, Ian Walmsley, and Frank K. Wilhelm, arXiv, December 2017 (30 pages).
- [1324] [Digital quantum computation with superconducting qubits](#), by Johannes Heinsoo, ETH Zurich, 2019 (271 pages).
- [1325] [Fast, accurate, high-resolution simulation of large-scale Fermi-Hubbard models on a digital quantum processor](#), by Gavin S. Hartnett, Khadijeh Sona Najafi, Aleksei Khindanov, Haoran Liao, Michael Schutzman, Michael R. Hush, Michael J. Biercuk, and Yuval Baum, arXiv, May 2026 (35 pages).
- [1326] [Evaluating the performance of quantum processing units at large width and depth](#), by J. A. Montanez-Barrera, Kristel Michielsen, and David E. Bernal Neira, arXiv, May 2025 (19 pages).
- [1327] [Quantum computers will not be that different: A blueprint for quantum computer architecture at scale](#), by Torsten Hoeffler and Matthias Troyer, arXiv, September 2026.
- [1328] QEC error syndrome detection costing more than the rest of the quantum computer comes from the first, unpublished, resource estimates from a couple vendors in the silicon spin realm.

This **System architecture** subsection contains 92 bibliographical references and notes containing at least 2,184 pages.

5.2 Error handling

One of the pitfalls of existing physical qubits is their significant error rates that are generated throughout the whole computing cycle, including qubit preparation, quantum gates and qubit readout and coming mostly from the fateful quantum decoherence as well as from defects from control electronics and photonics signals. Decoherence is mostly generated by the interactions between the qubits and their environment. It progressively destroys the quantum information sitting in the qubits, meaning superposition and entanglement. It leads to an inevitable failure in computation after a short time. Error rates depend on the operation, the device, the register size and the benchmark protocol. As of 2026, per-operation gate and readout error rates commonly sit between 0.1% and a few percent, depending on the qubit type and quality. But even 0.1% is an intolerable level for most calculations.

In conventional computing, errors are way less frequent but must still be corrected. While some errors may be detected and fixed during computing in microprocessors, most errors are happening in memory, storage and telecommunications. These errors are discrete, corresponding to some unwanted classical bit flips. With quantum technologies, errors happen first and foremost during computing and within qubits, and they are continuous and analog by nature, involving phase errors on top of the flip error that is akin to classical bit flip errors.

Errors are corrected through various mechanisms that we will cover here, without necessarily going into their details. This field is quite broad and is one of the most technical and cryptic you will ever find in quantum computing.

This subsection described first the source of errors and how they are measured. It then describes how errors are handled in NISQ (Noisy Intermediate Scale Quantum) systems with Quantum Error Mitigation and Quantum Error Detection and Suppression techniques.

We then describe the basics of quantum error correction, the various error correction codes available, including their related error syndrome decoding techniques, and then switch to fault-tolerance which is about supporting the execution of a large circuit fault-tolerantly. This is where we'll explain how logical operations are executed on logical qubits.

5.2.1 Error types and sources

We can semantically organize quantum computing errors with gate-based systems in three categories as presented in Figure 245, with “when” are errors happening during computation, what is their effect on the data, and where corresponding to their various sources at the classical and quantum physical levels.

When corresponds to the events during which these errors are generated during the quantum circuit lifecycle, starting with qubit reset, qubit gates, idle state and qubit readout. These correspond to typical error metrics like reset [1329], single and two qubit gates and readout error rates or percentages, or fidelities which are simply equal to 1 minus the error rate. In benchmarks that are used particularly with trapped ion qubits, SPAM (state preparation and measurement) errors measure the cumulative effect of qubit preparation and measurement. It is used and advertised by Quantinuum and IonQ. It however does not provide any indication on multiple qubit gates and qubits entanglement quality and scaling [1330]. All these error types are not qubit dependent and can happen with every qubit type, even with photonic qubits.

What is the effect of these errors on the data in a qubit register. Reviewing the literature drives a lot of confusion because the same words designate different objects, so we fix one vocabulary here, before discussing physical causes. A coherent error is a unitary over-rotation acting as $\rho \rightarrow U\rho U^\dagger$. A stochastic Pauli error applies a Pauli operator with some probability, as $\rho \rightarrow (1-p)\rho + pP\rho P$. Pure dephasing is the decay of the off-diagonal terms of the density matrix in the chosen energy or computational basis. Amplitude damping transfers excited-state population to the ground state and is nonunitary. Depolarization contracts the state toward $I/2$. Leakage takes the physical object out of its computational subspace and loss removes it altogether. An erasure is a leakage or loss event whose location is known to the decoder. Error correction folks work in the Pauli basis $\{I, X, Y, Z\}$ because syndrome extraction discretizes any in-subspace error into these components, which is why codes are designed to correct X and Z errors and their combinations.

Coherence times measure how long the signal of a specific decay experiment survives, and not how long superposition or entanglement lasts. Longitudinal relaxation is characterized by T_1 , transverse coherence by T_2 , measured with a Hahn echo, and by T_2^* , measured with a Ramsey experiment and including inhomogeneous broadening, these two being sometimes confused with each other in the literature [1331]. Under a Markovian model, $1/T_2 = 1/(2T_1) + 1/T_\phi$, T_ϕ being the pure dephasing time, a relation that holds only under that model. All these are fitted envelope constants [1332].

Errors are commonly described as “quantum channels” that are modifying the quantum state of the qubit which is represented by its density matrix ρ and how the error channels affect it.

Here is a segmentation of error types (“what”):

- **Coherent errors** correspond to a rotation in the Bloch sphere of an angle θ around a vector $\hat{n}$. It is a unitary and reversible transformation. This coherent noise is mostly due to a predictable control electronics imprecision. Its effect on the qubit density matrix





Figure 245: The when, what and where of errors affecting qubits separating logical operations when errors happen, their computational effects on the qubits and their various physical origins. In the “what” row, each entry is a distinct error model and they are not interchangeable. Leakage or loss becomes an erasure only when a detection mechanism locates the affected qubit without revealing the encoded information, and a maximally mixed state is unflagged depolarization, not an erasure. Qubit fidelities correspond to measurements done in the “when” realm, T_1/T_2 to the “what” realm, and specific benchmarks are done to identify the relative weights of the various physical error sources. The impact of physical sources of errors is variable and qubit type dependent. Another variation to be considered is the ability to dampen the impact of these errors or minimize them. For example, it seems possible to improve the quality of control signals and chips materials to avoid unwanted effects while it may be more difficult to avoid various many-body interactions or crosstalk effects. © Olivier Ezratty, 2023-2026.

T_1 : flip error (relaxation, dampening)

- qubit energy loss to the environment.
- time to decay from $|1\rangle$ to $|0\rangle$.
- spontaneous emission, quasiparticle tunneling, flux coupling, dielectric losses and control electronics imprecisions.
- more important at qubit readout.
- decay of qubit density matrix diagonal values.



decay is a continuous process

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}$$


T_ϕ : pure dephasing time

$$T_2^* \leq T_2 \leq 2T_1$$

$$Q = \omega_q T_1$$

Q = quality factor

ω_q = qubit resonant frequency



$T_2 - T_2^*$: phase error (dephasing, decoherence time)

- environment creates loss of phase memory.
- control electronics imprecision.
- important during computation.
- tied to number of consecutive gates.
- decay of qubit density matrix off-diagonal values.

Figure 246: Flip error and phase errors and their effect in the qubit Bloch sphere. © Olivier Ezratty, 2022.

Q How to differentiate errors, noise and decoherence?

These error related terms are sometimes used interchangeably but they have different meaning.

Errors is the generic term covering unwanted effects happening with qubits during initialization, gate operations and readouts which affect the quality of the results. The related metric are error rates and fidelities. Errors are generated by various sources of noise, but can also come from spontaneous emission, leakage like with superconducting qubits exiting their computational basis or with erasure errors, where the qubit is itself disappearing, which can happen with atoms and photons. Errors can also be generated by defects in qubit chips, or even by classical control and error correction software.

Noise is the main type of error source of stochastic nature that can be endogenous (like TLS errors or charge noise) or exogenous (quasiparticle poisoning, radiation effects, coherent errors coming from qubit drive signal imperfections, and thermal or electromagnetic fluctuations).

Decoherence is a particular effect of errors on qubit quantum states, generating dephasing, where the relative phase between quantum states becomes randomized, and amplitude damping, where the qubit relaxes to its ground state. It leads to progressively reduce superposition and entanglement in, and between qubits.

ρ is the channel: $\rho \rightarrow U\rho U^\dagger$ with $U = e^{-i\theta \hat{n} \cdot \vec{\sigma}/2}$, $\vec{\sigma} = (X, Y, Z)$ being the Pauli matrices. This type of error is reversible. It generates unwanted qubit rotations around the X ("bit flip"), Y and Z (phase flip) axis in the Bloch sphere. A small coherent over-rotation is not the same object as a Pauli operator applied with some probability, which is a stochastic channel $\rho \rightarrow (1-p)\rho + pP\rho P$ with $P \in \{X, Y, Z\}$. Coherent errors accumulate in amplitude over repeated gates while stochastic ones accumulate in probability, so they are budgeted differently. T_1 is

measured separately, by preparing $|1\rangle$ and fitting the decay of its population, the fitted exponential envelope reaching $1/e$ at $t = T_1$.

- **Amplitude damping** is one type of incoherent and irreversible error. It transfers population from the excited state to the ground state, which is a longitudinal transition, and damps the coherences at the same time. It is nonunitary and it is not a bit flip. It is described by two Kraus operators, E_0 with diagonal entries 1 and $\sqrt{1-\gamma}$, and E_1 whose only nonzero entry is $\sqrt{\gamma}$ in its upper-right corner, acting




Figure 247: How are measured T_1 , T_2 and T_2^* . © Olivier Ezratty, 2022-2023.

as $\rho \rightarrow E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger$, γ being the decay probability over the considered duration [1332]. The stochastic bit-flip channel $\rho \rightarrow (1-p)\rho + pX\rho X$ is a different channel and does not describe amplitude damping. These errors are related to the T_1 , which is linked to a loss of amplitude (“energy relaxation”), also called “longitudinal coherence time” or “spontaneous emission time”. It corresponds to a loss of energy in most qubits, coming from a dissipative process that releases some energy. In superconducting qubit circuits, T_1 is proportional to the circuit quality factor $Q = \omega_q T_1$, which itself is proportional to the ratio between the energy stored in the qubit resonator and its rate of energy loss. T_1 comes from different phenomena: spontaneous emission, quasiparticle tunneling, flux coupling and dielectric losses in the Josephson junction. The Purcell effect is a spontaneous emission through the readout cavity. The Purcell decay rate is related to the speed of this phenomenon. It is reduced using a Purcell filter which suppresses signal propagation at the qubit transition frequency. The filter is a pass-through with the readout cavity frequency that protects the qubit from decoherence channels while enabling its readout. Amplitude loss and coherence losses can always be decomposed in a linear combination of Pauli X and Z errors and identity. It explains why error corrections code’s main task is to correct Pauli X and Z errors. But error correction codes can be specialized in one type of error, like longitudinal errors [1333].

- **Coherence loss errors** happens when superposition is getting lost in the qubit. It contains depolarizing noise or pure dephasing which is related to the decay of the non-diagonal part of the qubit density matrix, creating a mixed state, shown up visually with the qubit state vector moving inside the Bloch sphere in an isotropic manner with the depolarizing noise channel being: $\rho \rightarrow (1 - \frac{3}{4}p)\rho + \frac{p}{4}X\rho X + \frac{p}{4}Y\rho Y + \frac{p}{4}Z\rho Z$. A maximally mixed state is complete depolarization

without any flag, and not an erasure: an erasure requires the decoder to know which qubit was affected. Another error of this type is pure dephasing, which is the decay of the off-diagonal terms of the density matrix in the chosen basis. It is neither a coherent Z rotation nor a variant of depolarization. Its effect on the qubit density matrix ρ is the quantum channel: $\rho \rightarrow (1-p)\rho + pZ\rho Z$. Without depolarization, these are phase flip errors. Longitudinal relaxation is T_1 , while T_2 and T_2^* both characterize transverse coherence, T_2^* including inhomogeneous broadening and the echo-derived T_2 refocusing part of the quasistatic and low-frequency noise. T_2^* is evaluated with a Ramsey experiment, applying one Hadamard gate, wait time t , then applying another Hadamard gate, and readout. Without phase errors, the probability should look like a sinusoidal curve. With it, the curve slowly converges around a probability 0.5. T_2^* is extracted by fitting the envelope of these fringes and taking the time at which the fitted envelope falls to $1/e$, and not by reading where a raw measured probability crosses $1/e$. T_2 is obtained with a Hahn echo experiment where an X gate is added at $t/2$, which removes some error sources not due to qubit defects (Figure 247).

- **Leakage errors** is another error type that sees a qubit drifting and stabilizing in another energy state than the basic $|0\rangle$ or $|1\rangle$, or when the qubit physically disappears like with ions and cold atoms. This can occur in the $|2\rangle$ level of a superconducting qubit, which we are trying to avoid, or with variations in the hyperfine energy levels of trapped ion qubits. Leakage errors are another type of erasure error. This type of error can be benchmarked [1334, 1335]. Two distinct families of methods address it. Pulse shaping such as DRAG (Derivative Removal by Adiabatic Gate) adjusts the control pulse to suppress the unwanted transition to $|2\rangle$ during driven gates, so it prevents gate-induced leakage rather than resetting

an already leaked qubit [1336–1338]. Leakage reduction units and reset or repumping protocols deal with leakage that has already occurred [1339, 1340]. It can use leakage reduction units (LRU) in error correction codes which are circuit based and next to a surface code QEC. These are also known as a leakage reduction circuit (LRCs) or leakage reduction units (LRUs), which use dedicated auxiliary qubits [1341, 1342]. Leakage error correction can also get the help of reading out the state of ancilla (or auxiliary) qubits that are processed as a qutrit, with the ability to measure three states instead of two [1343]. Leakage can also be generated by tunable couplers and requires yet other fixes [1344, 1345]. There is also an idle state information leakage that has a high degree of statistical significance and needs further attention [1346]. Photonic qubits do not leak into a higher energy level of the same object, but polarization encoding is not leakage-free: population escapes the encoded subspace into loss, multiphoton components and unwanted spectral, temporal or spatial modes [1347, 1348]. These play the same role in an error budget, with the difference that loss is usually heralded and enters fault-tolerance accounting as an erasure rather than a Pauli error [1349].

- **Stochastic Pauli noise** affects the qubit in a different manner across its dimensions. Its effect to the qubit density matrix ρ is the channel $\rho \rightarrow (1 - p_x - p_y - p_z) \rho + p_x X \rho X + p_y Y \rho Y + p_z Z \rho Z$, with different probabilities p_x , p_y and p_z affecting X, Y and Z.
- **Markovian and non-Markovian noise** is another classification of noise affecting qubits. Markovian noise is characterized by its memoryless nature, the noise at a given time depending only on the current state of the system and not on its past. It is described by a time-independent or time-local master equation, such as the Lindblad equation. The main Markovian noise errors are flip noise, dephasing noise and depolarizing noise. Non-Markovian noise has some memory effects with the qubit's evolution depending on its history. There are non-Markovian flip errors, dephasing errors, and noise errors with colored spectrum with different frequencies. It relates to $1/f$ noise in superconducting qubits that include charge noise and flux noise, $1/f$ charge noise with spin qubits and trapped ions $1/f$ noise coming from laser phase noise. Surface codes are sensitive to non-Markovian noise [1350]. Bosonic codes have better resilience to non-Markovian noise and specific error correction codes are also proposed to address it [1351].

One characteristic of qubit errors and their related T_1 , T_2 and T_2^* is their variability over time (it can change over time during computations or between computations, requiring frequent calibration) and space (it is qubit dependent) [1352]. In the quantum errors lingua, a **Pauli**

noise channel corresponds to a linear combination of flip, phase and depolarizing errors.

The goal of having long qubit energy relaxation times is in competition with that of achieving high-fidelity qubit control and measurement. One key concern is to be able to apply error corrections as fast as possible after they are detected and before qubit decoherence takes effect or gets amplified. This is the reason why readout gates and phase detection electronics (for superconducting/quantum dots spin qubits) must be as fast as possible. But fast gates and readouts can generate their own leakage errors! They also require high-bandwidth microwave pulses, which reduce the capacity to frequency multiplex it in readout microwave circuits.

In the literature, you will also find another noise segmentation with unital and nonunital noise. Unital noise doesn't favor any particular quantum state for example with depolarizing noise and phase flip noise whereas nonunital introduces a bias pushing the qubit toward a specific state, often the ground state. This is the case with amplitude damping that corresponds to some energy loss. This segmentation matters because unital and nonunital channels behave differently under a given code. How hard a channel is to correct depends on the code, the noise bias, the detectability of the event, leakage and loss flags and the hardware, rather than on unitality alone (Table 7).

	Unital	Nonunital
Unitary	Gate over-rotation, phase error	Empty by construction
Nonunitary	Depolarizing noise, phase flip, bit flip	Amplitude damping

Table 7: Classifying some noise sources in two domains: unital vs nonunital and unitary vs non unitary. Every unitary channel is unital, so the unitary and nonunital cell is empty as a mathematical consequence and not as a practical impossibility. Conversely, a nonunital channel is necessarily nonunitary. Correctability is model- and code-dependent.

Where are these errors coming from at the physical level. It can be various electromagnetic and thermal perturbations but various materials defects in the qubit chip circuits. It depends on the qubit type [1353]. At a very low scale, some qubit types are also damaged by cosmic rays, quantum vacuum fluctuations and gravity. These error sources are multiple [1354], leading notably to the progressive decoherence of qubits which affects qubits superposition and entanglement. They are linked to the various interactions between qubits and their immediate environment. Any operation will generate an error. An error can be generated at the time of correction, at the time of detection or at the time of application of a gate. Doing nothing on a qubit can also generate errors because of its finite coherence time.




Figure 248: sources of decoherence and cosmic radiations affecting superconducting qubits. Source: [1355].

Error sources include:

- **Control electronics** imperfections like clock, phase and amplitude jitter. It can be triggered by calibration errors of quantum gates that occur in the calibration of superconducting qubits. They can notably trigger leakage errors. These small errors can create imprecisions with qubit operations. They are mitigated by improving the precision of control electronics. With superconducting and quantum dots spin qubits, it is mainly located at the local oscillators, mixers, AWG (arbitrary wave generation) and DAC levels (digital to analog conversion). These are coherent errors while other errors drive qubits decoherence.
- **Crosstalk** describes situations where a given qubit may affect other qubits, usually through various electromagnetic fields transmissions [1357]. One typical crosstalk effect comes from using the same microwave frequency to drive distant qubits in a chip due to the bandwidth limitations of the microwaves that can be used for this drive.
- **Thermal noise** from components around the qubits. This is one reason for the existence of attenuators around superconducting qubits. It comes among other things from shocks between atoms.
- **Material defects** which are commonplace with solid state qubits (superconducting, quantum dots spins, NV centers). One common defect comes from dielectric losses in the Josephson junction of superconducting qubits [1358, 1359].
- **Electrical and magnetic noise**, which can have many origins depending on the qubits. It explains why quantum chips are integrated in some multi-layer metal enclosures to limit the impact of terrestrial magnetism on its qubits. Most solid state quantum processors are packaged in tight metal shielding.
- **Radiations** can be X-rays, gamma rays (whose electromagnetic nature was discovered in 1914), beta par-




Figure 249: How cosmic rays impact a Josephson junction. Source: [1356].

ticles and their ionizing effects as well as muons [1360–1363]. Radiation induces so-called multi-bit burst errors (MBBEs). Its sources are threefold: cosmic-ray-induced radiation (fix: underground location), terrestrial gamma rays in the laboratory environment (fix: surrounding the cryostat with a lead radiation shield) and naturally radioactive isotopes within the cryostat (fix: use radiopure alternative elements and use of an internal gamma-rays shield). The phenomenon is now well characterized, although only with small-scale chips [1364–1367], even for muons [1368]. Some of these radiations can be artificially generated by particle accelerators [1369]. In multiple quantum chips (superconducting qubits, silicon spins, quantum dots, etc.), these radiations create phonon quasiparticles in the chip substrate that can propagate to many surrounding qubits, endangering the efficiency of quantum error correction codes (Figure 249) [1370–1372]. Radioactivity is one of the many sources of long range correlated noise or correlated errors, also named error bursts [1373]. Envisioned solutions are to operate the QPU underground, which is not very practical [1374–1378], to shield the qubit processing unit with lead [1356] or copper back-side metallization [1379], to trap phonons using high impedance resonators made of granular aluminum [1380], adapting surface codes [1381], implementing distributed error corrections schemes [1382], gap engineering in qubit arrays, as proposed and experimented by Google AI in 2024 [1383–1385], using multiple chips to limit the propagation of phonons and quasiparticles [1386], with specifically engineered Purcell filters [1387], selective deactivation of qubits used in magic state factories to correct non-Clifford qubit gates [1388] and develop full-stack FTQC architectures like with the *Q3DE* model proposed in Japan [1389].

- **Mechanical vibrations** are also sources of correlated error in some qubit types like superconducting qubits. They can come from ambient noise, which

can be mitigated, as well as from the pulse tube (PT) in a dry dilution refrigerator [1390].

- **Vacuum quantum fluctuations**, which enter only through a specific coupling such as spontaneous emission into the electromagnetic environment, and whose quantified contribution is orders of magnitude below the noise sources listed above [1391].
- **Vacuum quantum fluctuations** are an endogenous source of errors within qubits while all the others are exogenous [1392]. They act through a specific coupling and their quantified weight in an error budget is negligible compared with the sources above.
- **Gravity**, mentioned here for completeness only. No quantitative model gives it a measurable weight in a qubit error budget, so it should not be listed on the same footing as crosstalk, materials defects or radiation [1393]. It could potentially be used to implement accurate gravity sensing using superconducting qubits [1394].

Generally speaking, errors are generated by various interactions, electromagnetic or mechanical, between qubits and their immediate environment and are associated with the phenomenon of quantum decoherence. Some of these effects are avoided by cooling the qubits to a temperature close to absolute zero, but this is not enough. Researchers are therefore working hard to ensure that the noise affecting the qubits is as low as possible so that qubits coherence time can be as long as possible.

We have to overcome this contradiction: qubits remain coherent, in a state of superposition and entanglement, if we do not disturb them, but we spend our time disturbing them with quantum gates operations! There are several ways to address these issues: improving qubit operations fidelity (preparation, gates, readout), improving the ratio between gate times and qubit coherence times, implementing quantum error mitigation techniques and quantum correction codes and at last, reduce the number of gates needed to run algorithms with various optimization techniques.

5.2.2 Qubits fidelities

In a gate-based quantum computer, three types of errors (or related fidelities) are usually evaluated: errors in single-qubit quantum gates, errors in two-qubit gates, and qubits readout errors. These error rates currently sit between 0.1% and several percent, depending on the operation, the device and the measurement protocol. Sometimes, idle qubit errors are also indicated, over a reference time inline with average gate duration.

In classical computing, errors are very rare. There are single particle perturbations (PPI) and single event upsets (SEU) which trigger “soft errors” or logical errors. The SER (Soft Error Rate) combines the SDC (Silent Data Corruption, not detected) and the DUE

(Detected and Unrecoverable Error, detected but not correctable). The unit of error measurement is the FIT (Failure in Time), which corresponds to one error per billion hours of use. The MTBF of electronic equipment (Mean Time Between Failure) is generally measured in years or decades. Errors are generally caused by isolated particles (ions, electrons, photons), particularly from cosmic rays like high-energy gamma rays. This affects in particular the electronics used in aerospace, which must be hardened to withstand them, as well as those used on Earth but at altitude. Memory is often more affected than processors. Thus, the error correction systems that use for example a parity bit and cyclic redundancy check used in telecommunications.

Qubits “fidelity” for any of these three dimensions is 100% minus the related error rate. In typical quantum parlance, when a “three-nines” 2 qubit-fidelity is mentioned, it means that it is better than 99.9%. In scientific literature, qubit fidelities are frequently presented with the formatting “99.914(3)%”, the number in parentheses describing the error margin of the last digit.

Figure 252 consolidates a comparison of some fidelity levels of superconducting, trapped ion and cold atom quantum computer qubits, this information being provided by their vendors. Qubit fidelities encompass these three fidelities/errors dimensions (1Q, 2Q, readout). You can find evolutions over time of these figures of merit in a paper from François-Marie Le Régent from Pasqal published in 2025 (coherence times, entanglement error, physical qubit count, experiments count, QEC code distance) [1395].

Two-qubit gates and readout error rates are generally higher than one-qubit gates error rates [1396, 1397]. We must therefore always pay attention to two-qubit gate errors, particularly given these gates are the source of much of the quantum computing power, through entanglement.

But these fidelities are not always measured in the same conditions. Some are measured with only a couple interacting qubits in research settings while others are done with all the register’s qubits being active and with tens of qubits. It can create significant differences favoring the first kind of measurements, particularly due to the significant crosstalk between qubits. As a result, fidelities are usually much better in research lab experiments than with commercial quantum computers. Thus, the choice to use only commercial settings in Figure 252.

Some best-in-class two-qubit gate fidelities achieved in commercial systems as of January 2026 were 99.92% for Quantinuum’s Helios 98-qubit trapped ions system, and 99.85% for IBM’s Heron r3 Pittsburgh 156-qubit superconducting system. The two-qubit CZ gates fidelity of China’s Zuchongzhi 3.2 107 qubits superconducting processor was 99.62% for an iSWAP [1398].

For all vendors, these error rates are currently prohibitive when executing many quantum gates in a typical

Q Is an idle qubit really idle?

The acceptable physical qubit error rate for a given quantum circuit, as shown in Figure 255, is roughly one divided by the circuit size which is its breadth in number of qubits times its depth in gate cycles. This accounts for idle operations, when the circuit contains no gate operation on given qubits. Physicists often label this an “identity operation”. But what is the physical hardware implementation of this operation? Is the qubit really physically idle? Not really!

In most architectures, matter qubits are continuously managed by electromagnetic fields:

- **Superconducting qubits** may receive microwave pulses for dynamical decoupling, which echoes out environmental noise. For flux-tunable qubits (e.g., from Google and Rigetti), a constant DC current conveyed by a “flux line” creates a static magnetic bias to park the qubit at a specific frequency.
- **Silicon spin qubits** are confined in semiconductor quantum dots using constant DC electric voltages applied to gate electrodes, while a static external magnetic field maintains the energy splitting of the electron spins. Even when technically idle, they often undergo continuous microwave pulse sequences for dynamical decoupling to actively cancel out ambient charge noise and residual nuclear spin fluctuations.
- **Neutral atom qubits** are optically confined by continuous laser light structured by SLMs and AODs photonic chips. This trapping light is often briefly turned off only when specific operations, like Rydberg excitation gates, are performed.
- **Trapped ions** are continuously confined by radio-frequency electric fields to stay in place within the trap. During long idle times, they may also undergo optical cooling, or be sympathetically cooled via neighboring ancilla ions.
- **Photonic qubits** are “flying qubits” that don’t stand still. An idle operation means it is passively propagating through a delay line such as an optical fiber or a circuit waveguide waiting to synchronize with other photons for a multi-qubit gate. In that situation, the qubit is inactive, but vulnerable to photon loss.

NISQ circuit. With each operation, errors accumulate and the success probability decreases, although not as a simple sum. Coherent, stochastic, correlated and leakage errors accumulate differently, and an error-corrected circuit follows other scaling laws. Imagine chaining a few dozen two-qubit gates. Under an independent per-operation error model, at a rate of 1% error per operation, the failure probability can very quickly exceed 50% at the end of a rather simple algorithm circuit and, generally, well before the fateful qubit coherence time limit.

Hence the fact that the power of a quantum computer is always evaluated not simply by the number of available qubits but by the number of operations that can be executed with a reasonable error rate at the end of the circuit run. To avoid this quantitative constraint, large circuits require error rates per operation many orders of magnitude below today’s physical values. There is no universal target. The desired error rates come from a circuit resource estimate with its depth, breadth and gate type decomposition, the target success probability, and a hardware noise model. Depending on the circuit size and hardware capacity, you can use an existing NISQ system or will need a future FTQC system. The discrepancy between today’s physical qubits fidelities and the need to run reliable large circuits explains the need for quantum error correction and fault-tolerance. At the edge of error correction exists a wall of time.

Reaching logical error rates per operation in the 10^{-16} range would allow even larger and longer circuits, but the corresponding execution times may become impractical.

A rule of thumb relates the gate error rate (e), the number of qubits (n) and the circuit depth (d). Under an independent per-location error model and for small error rates, the expected number of faults is about $n \cdot d \cdot e$, so a circuit with $nd < 1/e$ locations has a fair chance of running without any fault [1399]. This is a heuristic and not a definition of usable depth, which depends on the complete noise and circuit model and on the target success probability. As the error rate decreases, the usable circuit depth increases, and the range of usable algorithms expands (Figure 251).

For a quantum computer to be useful and scalable, you need a lot of qubits, a low error rate for quantum gates and qubits readout, and a long qubit coherence time to be able to execute algorithms without much time constraints although quantum error correction codes are indispensable workarounds for this last constraint.

How are quantum gates and measurement error rates evaluated (Figure 253)? We’ve seen previously how individual qubits flip and phase error rates are usually measured. Other methods are required to have an idea of the fidelities of registers with entangled qubits.


Circuit depth in the Noisy Intermediate-Scale Quantum Technology Era

Noise sets a limit on the maximum size of a computation without error correction.

Rough estimate:

$$nd \ll 1/\epsilon$$

n = number of qubits (width)
 d = circuit depth
 ϵ = error rate



Deep circuits → few qubits → efficient classical simulation.

Shallow circuits → many qubits → potential for a quantum advantage.

Figure 250: The empirical description of the acceptable errors to support a circuit with a given depth and breadth. Source: [1399].

Shallow circuits and their potential




Figure 251: Relationship between circuit depth and their use case. Source: [1399].




Figure 252: Scatter plot with average two-qubit gate fidelities and qubit number of commercial gate-based quantum computing systems. It highlights the most recent progress accomplished since November 2025 with, noticeably Quantum's Helios 98-qubit quantum computer released in November 2025 [1400], IBM with its 156-qubit Heron r3 generation of chips that nears the “three nines” threshold as well as its new Nighthaw 120-qubit chip, and QuEra systems. In some case like with the superconducting qubit systems from China, only a subset of the chip nominal qubits are usable due to defects on the chips and too low fidelities for either single qubit gates, two-qubit gates and readout. Data source: quantum hardware vendors two-qubit gate fidelities data obtained with randomized benchmarking (RB, although vendors may use different RB variants, and average values hide the tails, crosstalk and leakage that matter most for QEC) and plotted over QPUs number of qubits in log scales. © Olivier Ezratty, first version in 2023, revised in September 2026.

One method is the **Randomized Benchmarking** (RB) process which chains a random sequence of quantum gates whose result is known in advance and with comparing the results obtained with the right responses. Usually, a random sequence of Clifford gates is launched

and then executed backwards. The error rate increases with the number of chained quantum gates and depends on their type. We can evaluate the error rate of a given gate with the **Interleaved RB** which injects the gate periodically into the random gate set used. We

then measure the difference in error rate between the sequence with and without these added quantum gates [1401–1405]. You will have to look elsewhere to find out more data [1406, 1407]. The RB method has some drawbacks for clean noise quantification. It is apparently not suitable for the detection of every noise patterns [1408]. And there are many distinct protocols in this family, interleaved RB, simultaneous RB and cycle benchmarking among them, each with its own assumptions.

Quantum State Tomography (QST) is another method that we already covered in the section dedicated to measurement. It is based on a repeated measurement of qubit states which allows the reconstruction of a quantum system density matrix and the associated errors, for one or two qubits after some computation.

Yet another method exists that is based on some mathematical tools identifying a match between the noise rate of one and two-qubit gates of an algorithm and the total noise rate of the complete algorithm. In short, it links macro noise (algorithm) to micro noise (quantum gates).

Finally, **cross-entropy benchmarking (XEB)** helps verify large non-Clifford circuits such as those used in quantum supremacy experiments at Google. It runs instances drawn from an ensemble of random circuits and compares the measured output probabilities with the ideal ones computed classically, which is only possible in the classically verifiable regime. The resulting metric estimates a global circuit fidelity. XEB is sensitive to global circuit fidelity rather than average single-gate error rates.

5.2.3 Qubit figures of merit

After having described some key qubit figures of merit, let's reconcile all of them and see how they are connected to algorithms requirements in Figure 254 and Figure 255.

The qubit number relates to the circuit breadth. Qubit connectivity and the ability to execute multiple gates simultaneously conditions the code efficiency after its compilation.

Qubits relaxation, dephasing and depolarizing times condition the circuit depth, which is also the number of gate cycles in the algorithm. This is different from the gate fidelities which condition the circuit size, computed as the product of the circuit breadth (the number of qubits) and the circuit depth. The two-qubit gates used as a reference here are themselves dependent on the T_1 , T_2 and T_2^* of the two related qubits, plus some external error rates coming from control signals imprecision (electronics signal jitter, laser phase noise, etc.).

Two fidelities related figures of merit are now commonplace for physical or logical qubits, for the creation of Bell states and GHZ states. While a two-qubit fidelity characterizes only the gate, a Bell state requires at least an H and a CNOT gate, but also qubit initialization

and readout. Thus, Bell state fidelities are usually lower than for CNOT gates. GHZ state is an extension of such entangled state to 3 or more qubits, usually 4 in benchmarks. These characterize the ability of the system to create larger entangled states. It is becoming an important figure of merit as quantum computers start to scale.

Then, the case of photons is a bit special with specific figures of merit that we will cover later in the quantum hardware part. These are efficiencies, for photon sources, circuits and detectors, on top of their related error rates. This is generated by their probabilistic behavior. This has an impact on circuit execution times since any error requires repeating the related process, using various techniques like post-selection.

5.2.4 Quantum error mitigation

Quantum error mitigation (QEM) is about reducing quantum algorithms errors with combining classical post-processing techniques with some potential circuits modifications on top of running the algorithm several times and averaging its results (*aka the “expectation values of an observable”*, meaning the average of the observable eigenvalues weighted by the measured outcome probabilities).

QEM is a NISQ-era solution aiming at creating a quantum computing advantage before FTQC shows up in the longer term [1410, 1411] and for the support of deeper circuits. For example, in the case of IBM quantum processors, QEM enabled in 2026 the execution of circuits with about 12K two-qubit gate operations [1325]. A raw gate count is not a measure of circuit capability and should be read together with the circuit family, the observable, the target accuracy, the sampling overhead, the wall clock time and the unmitigated and classical baselines. Then, QEC and logical qubit will replace QEM to extend the reach of quantum computing. QEM reduces the influence of quantum errors using multiple runs and subsequent measurements coupled to some classical processing as opposed to QEC-based active qubits measurement and fast feedback-based corrections impacting the results of individual runs. QEM proposals started to appear around 2016 [1412–1414]. They do not share a single mechanism: ZNE extrapolates in the noise strength, PEC inverts a characterized noise channel at a sampling cost, symmetry verification and virtual distillation exploit structure in the state, and TEM builds a tensor network model of the noise. Only some of them learn a noise model from data [1415]. It is adapted to rather shallow circuits [1416] and we are not really sure whether it brings a quantum advantage on useful problems. Most QEM methods do not increase the required qubits count for a given algorithm. Many contributions in the QEM space come from IBM Research [1417] and Google AI.

Here are some identified QEM techniques:

Strategies for Characterizing Noise




Figure 253: Comparing the various strategies to characterize qubit noise. Source: [1408].


numbers	qubits #	}	circuit breadth
	connectivity: nearest neighbour 1x4, 1x2, 1x3, 1xN, N-to-N simultaneous gates = or < # of qubits		code and QEC efficiency
fidelities	SPAM (state preparation and measurement fidelity) reset/initialization fidelity: >99.9% readout fidelity: >99% for QEC leakage error rate (L): <0.1%	} harder ↓	circuit size
	single qubit fidelity: >99.9% two qubit fidelity: >99.9% Bell state fidelity: >99% GHZ n-qubit fidelity: >95%		
times	T_1 : relaxation time T_2 and T_2^* : dephasing time 1Q & 2Q gate time: conditions algo duration readout time: must be fast to enable QEC	} $\frac{\min(T_1, T_2)}{avg(gate\ time)}$	circuit depth
			circuit computing time QEC efficiency
efficiencies	photon source efficiency photon losses in nanophotonic circuits photon detection efficiency	}	circuit computing time

Figure 254: Consolidation of qubit figures of merit with their numbers, fidelities and time dimensions. On a practical basis, error rates / operation fidelities matter more than coherence times. T_1 , the Ramsey T_2^* and the echo T_2 are fitted decay constants of specific experiments, the 1/e convention applying to the fitted signal envelope. They do not translate into a 36% probability of obtaining a correct algorithm result [1409]. © Olivier Ezratty, August 2024-August 2025.

Zero noise extrapolation (ZNE) builds error models based on solving linear equations. It supposes that the noise is stable and cancels noise perturbations by an application of Richardson's deferred approach to the limit and works on short-depth (or shallow) circuits [1418–1420]. A similar protocol was designed for (hybrid) variational quantum simulators used to simulate the dynamics of quantum systems [1421]. Some ZNE methods are splitting decoherence error mitigation from depolarizing error mitigation [1422, 1423]. ZNE can even be adapted to photonic quantum computers [1424], to

GKP qubits [1425], and has been tested in combination with quantum error correction [1426].

Probabilistic error cancellation (PEC) is about detecting circuit bias with finding noise quantum channels, represented as density matrices for quantum gates, using quasi-probability decomposition of noise [1427–1430]. There's a sampling overhead in the process, which can be optimized [1431]. It then inverts a well-characterized noise channel to produce noise-free estimates of the algorithm observables [1418]. It can be implemented in mid-circuit measurement operations [1432]. Bayesian



D is constrained by the qubit coherence times, characterized by their relaxation and dephasing times T_1, T_2 and T_2^*

$N \times D$ is constrained by all error rates

$$\epsilon_{operations} < \frac{1}{N \times D}$$



corresponds to all qubit operations:

- single qubit gates errors
- two-qubit gates errors
- idle state errors
- readout errors

is conditioned by the ratio between T_1, T_2 and qubit operation times and by control originated errors

Figure 255: View showing the meaning of circuit breadth (number of qubits, rectangle height) and depth (gate cycles, rectangle width). The circuit depth is enabled by long relaxation (T_1), dephasing and depolarizing times (T_2) while the circuit size is conditioned by two-qubit gate fidelities which are themselves depending on the T_1, T_2 and T_2^* of their respective qubits. © Olivier Ezratty, September 2024.

readout error mitigation (BREM) corrects measurement assignment errors with a Bayesian model and is a distinct method based on different assumptions [1433]. Variations of PEC can handle non-stationary noise [1434]. It can also help reduce residual bit-flip errors from cat-qubits, enabling QEC to focus mainly on correcting phase errors below the 10^{-6} logical error rate [1435]. ZNE and PEC can also correct non-Clifford single qubit gates implementing arbitrary angle rotations, as used in NISQ variational algorithms ansatzes [1436]. The technique can be combined with error detection [1437].

Machine Learning Based Methods QEM are based on using training data to learn the effect of quantum noise in various situations [1438]. It is proposed by companies like QuantroOx (UK), by the University of Erlangen in Germany [1439], Quantum Machines (Israel). One of these is Clifford circuit learning or Clifford Data Regression is a variation of the previous technique that learns the effect of noise from Clifford gates using data comparing quantum emulation on classical hardware and runs on quantum processors. It then uses rather simple linear regression techniques to correct errors in post-processing [1440, 1441]. It can also be applicable to fault-tolerant T gates [1442]. Another proposed method uses physics-inspired machine learning [1443], graph neural networks [1444] and deep learning [1445]. In 2026, a CNN and transformer-based method was proposed in Germany to implement QEM on mid-circuit measurement operations used in NISQ distributed quantum computing (DQC) [1446].

Tensor Network Error Mitigation (TEM) is proposed by Algorithmiq to run quantum chemical sim-

ulations, on IBM's Heron QPU. In their paper, they showcase better results than ZNE and PEC, including shorter times to solutions [1447, 1448]. You also have QuantumNAS, a noise adaptative search method [1449] as well as some specific deep reinforcement learning (RL) techniques to improve qubit control precision [1450, 1451], including a novel technique based on LSTM neural networks [1452]. IBM is also investigating it [1438].

Readout Error Mitigation techniques are focused on correcting errors happening at the end of circuit with qubit measurement [1453–1455], using machine learning methods [1456], some being tailored for variational quantum algorithms [1457] and optimization algorithms [1458].

Error suppression by derangement (ESD) which provides an exponential error suppression with increasing the qubit count by $n \geq 2$ but is still adapted to NISQ architecture and shallow circuits [1459].

Other methods include symmetry constraints verification, distillation using randomized benchmarking [1463, 1464], randomized compiling [1465], applying gates simulating the reverse effect of errors [1466], reducing depolarizing noise [1467], quantum verification and post-selection [1468], doubling quantum resources in time with echo verification or in space with virtual distillation [1469], virtual distillation with derangement operators [1470], focusing on Pauli gates error mitigation [1471], using matrix product operators [1472], classically simulating the propagation of a noise model through a circuit and inverting its effect classically [1473], and also mixing various QEM and QEC techniques [1474–1476] like error filtration [1477], the “mitigated magic




Figure 256: Charting the Q-CTRL improvements. Sources: [1460] and [1461].




Figure 257: Segmentation of various quantum error mitigation and suppression techniques. Source:[1462]. Added in 2023.

dilution” technique proposed in 2025 by Riverlane to generate small angle rotations single qubit gates, based on quantum error mitigation technique [1478], and a technique embedding QEC to feed ZNE-based QEM with noise parameters data [1479].

Case specific application for chemistry [1480–1482], quantum signal processing [1483, 1484], and more gen-

erally, on sampling-based algorithms like the quantum phase estimate [1485]. QEM for analog computing is also envisioned, for quantum annealing using zero-noise extrapolation [1486, 1487] and Hamiltonian simulators [1488] and also for handling qutrits based quantum computing [1489].

Segmenting, detailing and comparing all these various methods remains an open challenge [1490, 1491], particularly considering that these different techniques can be combined [1492]. See Figure 257 that provides a high-level overview of the various quantum error mitigation and quantum suppression techniques.

QEM scalability is still questioned since its cost scales superpolynomially with the number of qubits [1493–1496] and also due to the quantum readout error mitigation techniques used in most QEM methods, which can introduce systematic errors that grow exponentially with qubit count [1497]. Also, QEM doesn't work well for quantum simulations of continuous dynamics [1498] and it can have both a negative and positive impact on the trainability of variational quantum algorithms [1499]. In 2024, a US research team established a result of this type in a restricted setting, showing for specified noise models, mitigation protocols and observables that the regime where mitigation succeeds is also classically simulable [1500]. This is a conditional statement and not a general property of every QEM method. Generic QEM incurs a sampling overhead that often grows exponentially with the circuit volume.

In 2024, Google with the chemical company Covestro (Germany) did implement QEM to advance the range of feasible VQE quantum chemistry from 12 to a mere 20 qubits [1501]. They found that to scale VQE beyond the capabilities of classical computing, they would need a 50-qubit QPU with all two-qubit gates having an error rate $\leq 3 \times 10^{-4}$, meaning fidelities above 99.97%. This is not far from the 2024 record of 99.87% obtained with Quantinuum with 56 qubits. In October 2024, they extended this work to combine QEM and EFTQC (early stage FTQC) to run a QPE [1502]. This method reduces qubit needs significantly but increases computation time.

5.2.5 Quantum error suppression and detection

Quantum Error Suppression (QES) is used to prevent or reduce the occurrence of errors during circuit executions. It makes the QPU more robust against noise and other forms of interference. QES methods include dynamical decoupling which applies a sequence of pulses to qubits to average out the effects of environmental noise, effectively decoupling the qubits from their surroundings, error-avoiding codes which use specific subspaces of the quantum system's Hilbert space that are inherently protected from certain types of errors [1503, 1504], such as decoherence-free subspaces, physical engineering that improves the design and materials of quantum hardware to reduce intrinsic noise sources and for example enhance qubit coherence times and minimize cross-talk between qubits and at last, specific randomization techniques implemented by code compilers like for the QSP algorithm [1505]. Qubit circuit cooling techniques also contribute to reducing the effect of thermal noise. Quantum Error Suppression can

be used to improve the error reduction rate of surface codes in quantum error correction, with the technique of noise-aware decoding [1506].

Dynamic Decoupling is a specific breed of error suppression techniques. It improves superconducting coherence times with decoupling idle qubits from other qubits under certain conditions [1507–1510]. It can also reduce single idle qubit decoherence by periodically refreshing their quantum state using control π -pulse gates (X and Y, and then -X and -Y). It seems that under certain circumstances, it can generate a good quantum speedup for oracle-based algorithms [1511]. It is also the base of the Space Curve Quantum Control (SCQC) codes designed at Virginia Tech. They are suppressing errors caused by quasi-static noises on qubits and work particularly with spin qubits [1512–1514].

Error suppression is widely used on IBM quantum processors to extend the circuit size they support [1515], with up to 13,800 two-qubit gates as of mid-2026, a count published without a common success metric and which should be read with its circuit family, observable, target accuracy, sampling overhead and classical baseline [1325].

Error suppression is one key specialty of **Q-CTRL**. In 2022, it announced that its autonomous error correction scheme could boost algorithms performance by up to $\times 9,000$ [1460]. Looking at the details, this is achieved with back-end and front-end optimization compilation and some error mitigation techniques. When you read their chart, shown in Figure 256, you find that the $\times 9,000$ ratio pertains to the success rate of running a Bernstein-Vazirani algorithm on a superconducting qubits processor, in the case of 16 qubits. But the related success factor is below 20% and is an extreme case. You must remind yourself that 16 qubits can be easily emulated on your own laptop and is way below any quantum computing advantage. If you were to extend their chart beyond 30 qubits, you'd be hundreds of thousands better than their competitors but with a very small success probability.

Later, Q-CTRL used QEM to a gate-based QPU from IBM to perform better than D-Wave quantum annealers to solve larger combinatorial optimization problems with better chances to yield correct results [1516]. It drove some pushback from D-Wave and counter-pushback from Q-CTRL [1517]. Whatever, both solutions are worse than best-in-class CPLEX based classical solution.

Quantum Error Detection is an error correction mechanism that is proposed for early FTQC regimes where error syndrome measurements are executed at predetermined times and the related computation is discarded when an error is detected, generating an obvious computing time overhead [1518]. It leads to the so-called post-selection techniques which can't scale for real-size quantum circuits, although this effect can be somewhat mitigated [1519, 1520]. If each of 1,000 steps is accepted independently with probability 0.99, a single run sur-

vives all of them with probability $0.99^{1000} \approx 4.3 \times 10^{-5}$, so the mean number of runs before one is accepted is the reciprocal of that, about 23,164 runs. And we expect FTQC algorithms to become useful with at least a million operations! Error detecting codes have a distance $d = 2$ in the $[[n, k, d]]$ notation that we describe a little later in 5.2.6. A $[[4, 2, 2]]$ can only detect one type of error, usually, flip error, and not correct it. It is tested for small scale NISQ circuits like QAOA [1521].

Some experiments implement pseudo-logical qubits using quantum error detection and suppression, like in 2026 on an IBM heavy-hex quantum processor [1522]. It shouldn't however be assimilated with error correction and fault-tolerance.

Post-selection is still needed with FTQC platforms for specific features. It is used in the magic state distillation process and "T gates factory". Inside the factory, post-selection is used to throw away any "dirty" states. These factories are local and small and they can run in parallel with the circuit, preparing T states that are then used "on demand". Also, checking errors makes use of auxiliary qubits. But they can also generate some errors. One option is to perform a quick check and post-select only the auxiliary qubits that are confirmed to be in the correct $|0\rangle$ state before using them.

Quantum error detection, suppression and mitigation can be combined together to support larger circuits. See this long 2026 review on all three [1523].

5.2.6 Error correction basics

The basic operating principles of quantum error correction are illustrated in Figure 258 with an eight-step correction, using the example of a simple flip-error correction code [1396, 1524–1526]:

1. **Encoding:** the qubit $|\psi\rangle$ to be corrected will first be coupled a certain number of times via CNOT gates with several data qubits (here 2). The resulting qubits are entangled. In the example, we get the state $\alpha|000\rangle + \beta|111\rangle$ for an input state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. This is called a repetition code. The three entangled qubit become a "logical qubit" $|\psi_L\rangle$.
2. **Channel:** single computing, idle operation, or transmission operation that potentially generates an error coming from various sources of noise. The error source is applied to the three entangled qubits.
3. **Detection:** one or more error syndromes are detected via quantum gates that associate qubits with auxiliary qubits and using parity checking. The example in Figure 258 detects flip errors. The detection is usually based on stabilizer codes which check the parity of two or more qubits using H and CNOT or CZ gates. A change of basis with an H gate before this operation enables syndrome detection in the Z basis (when using a CNOT) or X basis (when using a CZ), for a phase error detection.

4. **Measurement:** the state of these parity checking auxiliary qubits is measured to generate classical bits. This is also called a syndrome extraction process. This must be some non-demolition measurement for the corrected entangled qubits that does not destroy it. These measurements are labeled "mid-circuit measurements" since it occurs before the end of your circuit execution, only on a subset of the register qubits and without changing the quantum state of other qubits in the register. The syndrome measurement part uses here two "ZZ measurements", the stabilizers Z_1Z_2 and Z_2Z_3 , which jointly measure the parity of qubit pairs in the Z basis and detect bit-flip errors without revealing the encoded state. The phase-flip repetition code is obtained by conjugating the encoding with Hadamard gates, and its stabilizers are then of X type, X_1X_2 and X_2X_3 .
5. **Decoding:** using the classical result of mid-circuit measurement, the syndrome detection takes place classically to determine the error to correct. In the flip error case, it generates the index of the qubit to be corrected in the replicated entangled qubits using a simple correspondence table. It cannot detect an error affecting simultaneously two or three of the entangled data qubits. With more complicated error correction codes like surface codes, this decoding process can be costly, but its complexity is not generically exponential. Practical decoder families, matching, union-find, belief propagation, tensor network and learned decoders, run in polynomial and often near-linear empirical time, the binding constraints being latency and accuracy rather than asymptotic scaling [1527]. Decode latency is indeed a critical figure of merit of QEC [1528].
6. **Correction:** the decoder output is used to recover the encoded state, either by applying a single X gate, or a Z gate for phase errors, or, more commonly in practice, by updating a classically tracked Pauli frame in classical control software without applying any physical gate. The decoder identifies an equivalence class of errors modulo stabilizers and not necessarily the exact physical error. This error correction code cannot correct an error that would affect more than one of the three entangled qubits. The parity measurement corrects small errors when they don't need a correction.
7. **Recycling:** the data qubits are left entangled to be reused for the next correction round. Decoding the block back into an unencoded qubit is not part of a normal correction round and would break fault tolerance if performed between rounds. The state stays encoded until the final logical readout. The two auxiliary qubits are reset to $|0\rangle$. Some QEC protocols avoid this reset under some circumstances [1529].
8. **Reuse:** the data qubit and the auxiliary qubits can be reused for subsequent operations that will also be corrected using the same process.

Error correction codes circuits such as Shor’s 9-qubit code on [Wikipedia](#) are usually not complete. They usually don’t show the measurement and correction steps. These charts do not explain where error correction codes are implemented when running a quantum circuit. The reason that we will detail later is that error correction codes correct a memory made of logical qubits to create virtual logical qubits with a longer coherent lifetime than their underlying physical qubits. Error correction is a continuous and repetitive process that stabilizes a memory made with logical qubits. This continuous process is interspersed with logical gates acting on the logical memory, with different implementations as we will detail in 5.2.11. Indeed, some logical gates in the Clifford group are not necessarily physically implemented in the logical memory. This process is completed by a separate magic state generation and their teleportation on logical qubits.

Two distinct tasks are involved here. Architecture selection and resource estimation choose the code, its distance, the magic state factories, the decoder and the physical layout, from the circuit size and the target logical error rate. The compiler then maps an already selected FTQC scheme, handling logical gate synthesis, placement, routing and scheduling, and may optimize it [1530, 1531]. It must also allocate resources for the magic state factory and teleportation. All this is circuit-dependent. QEC codes will increase computation time by a scale depending on the time overhead of error syndrome measurement and detection. Error correction must indeed be accounted for when evaluating the quantum computing speedup gain brought by a given algorithm with another one or with a classical equivalent.

Is the simple flip-error correction code in Figure 258 a three or five qubits code? It is usually described as a 3-qubit code when dealing with the size of the logical qubit $|\psi_L\rangle$ or a 5-qubit code when including auxiliary qubits used in syndrome detection (3+2=5).

A variation of this code adds some H gates to correct phase errors. The 1995 **Shor’s 9-qubit error correction code** is a simple code consolidating these two methods, with the corrected qubit being replicated 8 times, as shown in Figure 259. This code corrects both flip and phase errors. Its code block contains 9 data qubits. The figure of 15 qubits sometimes quoted corresponds to one particular syndrome extraction method using 6 ancilla qubits, the total hardware count depending on the extraction design, the verification ancillas and whether ancillas are reused. In the first phase, the corrected qubit is replicated two times, and each resulting qubit is again replicated two times with CNOT gates. The first three blocks of 3 qubits implement a flip error correction. It outputs 3 qubits which then implement a phase error correction.

Raymond Laflamme (1960, Canada) demonstrated in 1996 that at least five physical qubits are needed to create a “logical qubit” integrating flip and phase error

correction. With **Emanuel Knill**, he also demonstrated that any single qubit error was a linear combination of flip and phase errors, leading to factoring error correction into flip- and phase-error corrections [1532, 1533].

In practice, the 7-qubit **Steane** code is the most referenced because it is not redundant like the Shor code. These 3-, 5-, 7- and 9-qubit codes are part of a generic group called **stabilizer codes** formalized by **Daniel Gottesman** in 1997. We are now going to dig a little deeper into how they operate.

We will better understand how an error correction works without reading the state of the qubit to be corrected. Let’s take the case of a simple flip error correction code with three qubits.

These three entangled qubits can have an error X_1 , X_2 or X_3 or no error (I =identity). X is an amplitude inversion Pauli gate. It creates an amplitude inversion of the corresponding entangled qubit as shown in the equations below. These new states correspond to three errors and the absence of errors. These four states have the interest of being mathematically orthogonal for all the values of the α and β defining the state of the qubit to be corrected. The trick is to perform a measurement of these values in the vector space corresponding to these four values and not in the original qubit computational base. This will not deteriorate the coherence of the original data qubit.

$$\begin{aligned} |\psi_L\rangle &= \alpha|000\rangle + \beta|111\rangle \xrightarrow{I} \alpha|000\rangle + \beta|111\rangle \\ |\psi_L\rangle &= \alpha|000\rangle + \beta|111\rangle \xrightarrow{x_1} \alpha|100\rangle + \beta|011\rangle \\ |\psi_L\rangle &= \alpha|000\rangle + \beta|111\rangle \xrightarrow{x_2} \alpha|010\rangle + \beta|101\rangle \\ |\psi_L\rangle &= \alpha|000\rangle + \beta|111\rangle \xrightarrow{x_3} \alpha|001\rangle + \beta|110\rangle \end{aligned}$$

The syndrome extraction is called a “stabilizer code” or “stabilization code”, which will feed the ancilla (or auxiliary) qubits. The process is the same to evaluate and correct a phase error but with Z gates instead of X gates. The disadvantage of the solution is that it cannot detect errors that would occur simultaneously on two or three of the entangled qubits. No error-correcting code can correct all errors!

The stabilizer codes formalism generically describes the error correction codes we have just studied with three parameters describing a code block: $[[n, k, d]]$ with:

- **n** = number of physical data qubits used in the code, with $n > k$. These physical qubits store redundant information thanks to entanglement. This number, *aka* blocklength, doesn’t include additional auxiliary qubits.
- **k** = number of logical qubits. $k = 1$ for surface codes which spread data qubits in a code patch to implement a single logical qubit and $k > 1$ for qLDPC which merge multiple logical qubits on a given num-

3-qubit flip error correcting code in quantum computing




Figure 258: Simple flip-error correction code with its 8 steps. Idealized, non-fault-tolerant illustration: ancillas, gates and measurements are assumed perfect and the correction is applied directly on the data qubits. Adapted from [1524].




Figure 259: Full Shor 9 error correction code correcting both flip and phase errors. Idealized, non-fault-tolerant illustration: the nine qubits shown are the code block, while the ancillas used for syndrome extraction, their verification and the measurement and correction steps are not represented. Adapted from [1534].

ber of physical data qubits, with a better encoding rate *aka* “design rate” ($k/n < 1$).

- d = the weight of the smallest undetectable logical operator, meaning the smallest number of simultaneous

qubit errors that can transform one valid codeword into another, *aka* the **code distance**. In the ideal code-capacity picture, a code of distance d detects up to $d - 1$ errors and corrects up to $t = \lfloor (d-1)/2 \rfloor$ of them (1 for $d = 3$, 2 for $d = 5$ and 3 for $d = 7$). This counting rule is a property of the code and says nothing by itself about performance under circuit-level noise, which depends on the decoder, the syndrome extraction circuit and the noise model. This t number is the weight of the error to correct. In classical error correction, it corresponds to the Hamming distance. In the 2D representation of surface codes, the code distance corresponds to the dimension of the data qubit code square, which is 5 in Google Sycamore in 2023 and 7 in the 105-qubit Willow chip introduced in 2024 in Figure 274.

An error correction code with a code distance d can correct errors for up to $(d - 1)/2$ (replicated) qubits, or say differently $d \geq 2m + 1$, m being the number of redundant qubits that can be corrected. You need d to be at least 3 to correct both flip and phase errors.

In this notation, Shor’s 9 qubit code is a $[[9, 1, 3]]$, Steane’s is a $[[7, 1, 3]]$ and Laflamme’s smallest code is a $[[5, 1, 3]]$ (demonstrated in September 2026 by AQT for implementing a teleportation between two logical qubits [1535]). They all have a code distance of 3. A simple 3-qubit flip or phase correction code is a $[[3, 1, 1]]$ code: as a full quantum code its distance is 1, since it offers no protection against the error type it was not designed for. Against its targeted error model alone, bit flips for the first and phase flips for the second, it behaves like a distance-3 classical repetition code, which is why the restricted error set must always be stated. There are larger cases like with the $[[512, 174, 8]]$ CSS code [1536]. Named after their creators Calderbank, Shor and Steane, these are a family of quantum error-correcting codes combining two classical error-correcting codes, one to correct bit-flip errors and the other to correct phase-flip errors. Thanks to their linear combination, they can correct all Pauli errors. Beyond Shor-9 error correcting code, which is the simplest CSS code, the choice of stabilizers conditions the efficiency of the code. It can be optimized using machine learning [1537].

The ratio k/n is the code encoding rate, efficiency or design rate (<1), n/k is the code overhead and the number of physical qubits to build a logical qubit (>1), and d/n gives another indication on the code overhead to obtain protection from errors. These determine the quality and efficiency of error correction codes. For example, the Shor-9 flip and phase error correction code is a $[[9, 1, 3]]$ code with a low efficiency while many modern qLDPC codes aim for high rates where n is only slightly larger than k . With most codes, the total overhead including auxiliary qubits is about $2n/k$. In a $[[n, k, d]]$ code, a code block uses n physical data qubits to encode k logical qubits. One reason for qLDPC codes being appreciated is that they have a number of logical qubits

k and a code distance d that both scale linearly with the number of physical qubits n .

The **geometric topology and connectivity** can be regular 2D like with Google (for topological codes like surface codes), 2D/3D for color codes like the 7-qubit Steane code (which is a 2D code with distance $d=3$) and qLDPC codes for high or long-range connectivity (like the bivariate bicycle codes used by IBM). The qubit physical connectivity indeed influences the choice of error correction codes.

Some codes are also specialized to correct a specific type of noise. When your hardware, like cat-qubits, has way more phase-flips than bit-flips, you don’t use a low encoding rate 2D code like the surface code, but either a so-called XY-surface code or a qLDPC code. Codes can also be combined like qLDPC and concatenated codes [1538].

Finally, error correction codes decoding speed requires the use of classical computers that are responsible for reading and interpreting the errors and calculating the fix. It must cope with the gate operation speed and qubit coherence times. Surface codes are fast and local while qLDPC codes are more efficient in physical resources but harder to decode. They are therefore somewhat more adapted to “slow” qubit modalities like cold atoms and trapped ions.




Figure 261: Error correction replacing measurement with a controlled operation. Source: [1540].

The stabilizer codes use a syndrome table that provides a match between the errors on each qubit and the detected syndrome. The number of auxiliary qubits used to create this table must therefore be sufficient to identify the qubits to be corrected in the logical qubit. In the example in Figure 260 with a logical qubit with 7 physical qubits, the 3 auxiliary qubits allow the identification of eight scenarios, sufficient to determine which of the 7 physical qubits must be corrected. The 8th scenario is the absence of error, therefore needing no correction.

Some quantum error correction codes detect syndrome and correct errors without using measurement and classical syndrome detection steps, parity checks, and correction being directly implemented with quantum gates (see Figure 261 on the bottom). This are labeled a

7 qubits error correction code named $[[7,1,3]]$ in the stabilizers formalism




Figure 260: A 7-qubit correction code with a code distance 3. Source: [1539].

“no-ancilla-error-detection” codes (NAED) [1541]. It is not commonplace for a reason because it seems to require very high fidelities physical qubits with error rates of about 6×10^{-5} . It is envisioned with trapped ions and cold atoms [1542], and with 2D and 3D color codes [1543].

Measurement-free quantum error correction (MFQEC) is one such technique [1544]. It requires an unconditional qubit reset gate and it corrects only single gate errors. A full FTQC architecture using this technique was proposed in 2025 by AQT (trapped ions) [1545].

Then, you have the autonomous methods also branded **aQEC** (autonomous quantum error correction) that are implemented at the qubit level. Some energy dissipation must be handled, using a technology called reservoir engineering, which is implemented in cat-qubits and other bosonic codes [1546–1548]. Otherwise, whatever, the auxiliary qubits used in typical QEC are reset to $|0\rangle$ and this reset is a dissipative process. The thermal bath is just elsewhere!

Threshold define the higher bound of physical qubit error rates for a quantum error correcting code to operate successfully. It depends on the code itself and on the qubit type and on the system noise model. Concatenated codes have a low threshold while surface codes have a higher threshold, about 1% for a commonly used circuit-level depolarizing model with a matching decoder, and are thus tolerant to higher qubit error rates. That 1% is illustrative of one named model and not a universal constant. But they require a larger number of physical qubits per logical qubit. Physical qubits must be connected to their immediate neighbors in a 2D structure or with honeycomb variations [1549, 1550]. At threshold, increasing the code distance stops suppressing the logical error rate asymptotically, so no finite distance

reaches a given low target. This is what the divergence in Figure 265 expresses, rather than a physically infinite code. For example, if the physical error rate is 10 times below the threshold, the surface code distance d should be greater than 17 to achieve a logical error rate below 10^{-10} [1551]. How a code threshold computed on a given physical system is quite a complicated matter [1552].

Break-even is the situation when a logical qubit outperforms an explicitly stated physical baseline, usually a longer memory lifetime than the underlying physical qubits. It is a different criterion from below-threshold scaling, which is about the logical error rate decreasing as the code distance grows. You usually reach break-even when physical error rates are at least one order of magnitude below the threshold, since near threshold the distance required for a given logical target grows very quickly. To obtain a reasonable number, you need better qubit fidelities as shown in Figure 265.

Gadgets are specialized quantum circuits or sub-routines designed to execute necessary logical operations, such as syndrome measurements or non-transversal gates, on encoded data while preventing physical errors from cascading. Gadgets are engineered to contain fault propagation. They guarantee that a single component failure within the sub-circuit, such as an ancilla qubit error, translates to at most one physical error on the target data block, and preserve the integrity of the protected quantum information.

5.2.7 Error correction codes

Quantum error correction can’t work the same as classical error correction. Qubits cannot be independently replicated with some measurement that would be performed on one replicated qubit, per the famous no-cloning theorem. On top of that, we are correcting

analog errors in multiple dimensions, not just a 0/1 error flip that could be labeled as a simple “digital error” [1553–1555].

The techniques explored for over two decades is about implementing quantum error correction codes called **QEC** for Quantum Error Correction or rather **QECC** for QEC Codes. QEC uses redundancy through encoding quantum information into a larger Hilbert space using qubit entanglement.

QEC is about creating logical qubits which assemble multiple physical qubits with a number that depends on many factors: the error code, the target logical qubit fidelity, the algorithm size, the physical qubits quality and connectivity. QEC codes usually correct bit flips (X errors) and phase flips (Z errors), which makes it possible to correct any linear combination of these two errors. QEC detects and corrects these errors by mapping the state of the qubits to correct into a larger state space where the presence of errors can be detected without directly measuring the encoded quantum information. Syndrome measurements detect errors without collapsing the quantum state of the logical qubit. Then, (classical) syndrome decoding helps determine which qubit must be corrected and corresponding recovery operations are applied to correct the error and return the system to its original state.

Most of these QEC schemes correct errors that are small and independent, meaning, not correlated between several close and distant qubits [1556, 1557]. Specific error correction codes still exist to correct correlated errors like those which are generated by cosmic rays with solid state qubits [1558] and leakage errors when the qubit exits its computational space, like with superconducting qubits.

Error correction codes apply to both universal gate quantum computing and quantum telecommunications. In the first case, they are integrated into the broader concept of fault-tolerant quantum computing (FTQC). Error correction’s effect is to slow down qubits decoherence and to extend the available computation time. Figure 262 makes an inventory of the main quantum error correction codes with their origin and date of creation [1559]. This error correction zoo is very dense and continuously expanding [1539, 1560, 1561]. It includes several families of error correction codes. Their generic design requirements are to minimize the physical qubit error requirements (error threshold level), the number of physical qubits per logical qubit, their connectivity needs, and the time and classical computing resources to decode error syndromes.

In quantum computing, quantum error correction also has variations adapted to memory qubits like for a qRAM, which may have different constraints. In memory qubits, error models are predominantly due to decoherence, including bit-flip and phase-flip errors. It usually fault-tolerantly correct only Clifford group gates. It focuses on keeping the coherence of qubits over time.

Codes like the 9-qubit Shor code or surface codes are used to detect and correct errors periodically. They usually don’t need to implement costly magic state distillation since memory qubits don’t need to support a universal gate set. They keep track of data teleported from computing qubits that may have been prepared fault-tolerantly with a universal gate set. Qubits candidates for quantum memories include electron spins coupled with nucleus spins in NV-centers and other cavities [1562].

Stabilizer codes use a set of parity-check operators to correct flip and/or phase errors with five (Laflamme), seven (Steane) or nine qubits (Shor). These codes couple qubits several times using entanglement. Then, some quantum code using auxiliary qubits are used to detect error syndromes without affecting the corrected qubits. It checks whether pairs of qubits have the same value on the initial basis or using another basis after a rotation, to detect phase instead of flip errors, using so-called error syndrome measurements. All this is done without reading the value of the initial qubits which would make the whole system collapse. We then use single-qubit gates to correct the qubits for which an error was detected. It goes through some classical processing that must be as fast as possible that creates a correspondence between the error syndromes and the qubits to correct. In large surface codes, this process can be quite expensive [1563].

Notice the existence of **Crumble** which is a prototype stabilizer circuit editor that can be used to test stabilizer codes. Also, Peter Shor et al. proposed in 2025 to use a graph representation of stabilizer codes to improve their construction, using the ZX calculus graphic language (that we describe in the software tools part of this book) [1564].

Stabilizer codes have many variants including:

Quantum LDPC codes, or qLDPC codes, are QEC codes whose stabilizer checks are sparse in the same sense as classical LDPC codes [1565, 1566]. Sparsity means that each stabilizer acts on only a constant number of qubits and each qubit participates in only a constant number of stabilizers, even as the code distance grows. This structure makes syndrome extraction and decoding scalable. Sparsity alone does not guarantee a constant rate, a linear distance, a high threshold or an efficient decoder, which are all properties of particular families. Some qLDPC families reach high encoding rates, $1/2$ or $1/5$ being high in this context compared with the $1/(2d^2 - 1)$ of a surface code [1567], together with a growing distance, which is what makes them attractive for large logical circuits. The asymptotically good families constructed by Panteleev and Kalachev encode a number of logical qubits and reach a distance that both scale linearly with n [1568], while a surface code encodes a single logical qubit in about d^2 data qubits. Many finite qLDPC constructions require non-local *aka* long distance checks between qubits, an implementation constraint discussed below.




Figure 262: Inventory of key quantum error correction codes, the code in the yellow zones being mentioned in this part.
 © Olivier Ezratty, 2023-September 2026, inspired from [1559].

There are several families of qLDPC codes [1566]. Hypergraph product codes provided the first explicit construction with nonzero rate [1569], and later developments such as balanced product codes, lifted product codes, and the Panteleev-Kalachev “good” quantum LDPC codes improved distance scaling and decoding performance [1568]. These codes work on the condition of the quantum processor supporting long-range connectivity, such as with trapped ions, photonic networks, or superconducting systems with long-range connectivity like with IBM’s Loon architecture that will be implemented in future chips after the Nighthawk chip that went live in January 2026 with its 1-to-4 qubit connectivity.

qLDPC codes indeed inherit their core philosophy from classical LDPC codes, which were introduced in the 1960s and later rediscovered as some of the most powerful classical error-correcting codes [1565, 1568, 1570–1575]. In classical error correction, LDPC codes exhibit excellent performance because their parity-check matrices are sparse, enabling efficient iterative decoding on large graphs. qLDPC codes adopt this same structural idea with stabilizers remaining low-weight and each qubit participating in only a few checks. Another benefit from high-encoding rates qLDPC codes is their ability to support distributed quantum computing architectures [1576], including the Q-Fly architecture proposed by Rodney Van Meter et al. from Keio University in 2026 [1577].

qLDPC codes require a better qubit connectivity than surface and color codes. This can be based on using single or multilayer connectivity chips for superconducting qubits, as is being designed by IBM which initially studied surface codes [1578, 1579] and then switched to working with qLDPC when they realized they could implement long range connectivity with their qubits using only two layers of additional coupling wave guides in the Loon layout announced in 2025, with four nearest-neighbor connections and 2 distant connections for each physical qubit. IBM designed its own qLDPC code called the gross code [1580, 1581]. This variant of bivariate bicycle qLDPC code uses an $[[144, 12, 12]]$ encoding and the name gross comes from 144 being a “gross” (a dozen dozen). IBM later proposed a double-gross code $[[288, 12, 18]]$ supporting a larger distance with only dividing by two its encoding rate ($k/n = 12/288$ instead of $12/144$). Robert J. Schoelkopf (Yale, QCi) proposed a similar approach and qubit connectivity in 2026 which is adapted to dual-rail superconducting qubits [1582].

qLDPC codes were also successfully tested with cold atoms qubits using Rydberg state driven non local connectivity [1583] or with moving them around from a storage zone to a dedicated zone for the execution of Rydberg state-based two-qubit gates [1584]. They currently support FTQC logical operations mainly within the Clifford gates group. There are variants in syndrome decoding to improve their efficiency and time [1585–1588] and with special versions adapted to bias-noised error

qubits [1589], including bosonic qubits [1590] as well as variants with 1D and 2D qubit nearest-neighbor layouts [1591]. But new qLDPC variants seem to be able to implement transversal non-Clifford gates like CCZ and T [1592, 1593].

There are also some proposals to create hybrid qLDPC and surface code QEC [1594], and many variants like the Mirror codes which do not rely on (linear) CSS stabilizers and exhibit better fault tolerance [1595], the Self-dual Stacked qLDPC [1596], non-Abelian qLDPC codes [1597], spatially coupled codes that work well with low-latency decoders [1598], codes that can also correct efficiently Clifford group qubit gate errors [1599] and correct (future) quantum memories [1600]. All these codes have various sets of trade-offs and inherent mathematical complexity. Up to a point that some researchers in the field are experimenting the use of LLMs to design new qLDPC variants [1601].

Q Where is this bivariate bicycle naming coming from?

The “bicycle” naming was created by David MacKay, Ian Mitchison, and Paul McFadden in a 2003 paper introducing a family of quantum sparse-graph codes [1602]. In these codes, the parity-check matrix $H = [A \mid B]$ is built by concatenating two square circulant, or cyclic, matrices A and B . Since a circulant matrix can be visualized graphically as a cyclical ring of nodes, placing two of them side-by-side conceptually creates a structure with “two wheels”, so, a bicycle.

The “bivariate” naming has another origin, to make it simple, when IBM team replaced in 2024 the single cyclic variable used in the bicycle code with two variables [1603].

One caveat is that it is not possible to simultaneously have constant rate, growing distance and short-depth gadgets to perform encoded CNOT gates in such codes [1604]. As a result, the reduced physical qubits overhead of qLDPC codes is obtained at the expense of longer circuits, on top of more complicated (classical) correlated error syndrome decoding, due to errors affecting multiple logical qubits simultaneously. This is a costly global optimization problem. Belief propagation, usually completed with post-processing such as ordered statistics decoding, is the most common family for qLDPC codes, but it is not the generic answer to all correlated decoding, other codes relying on matching or other decoder families (see 5.2.8) and the challenge is to implement it in real time, particularly for large distance codes [1605, 1606].

There are also trade-offs between the number of required layers or connectivity requirements, and encoding rates gain. These are studied and compared by MIT [1607]

and also in China, in Guo-Ping Guo’s team and Origin Quantum, with their “routing codes” [1608]. Such improved connectivity is also investigated with silicon spin qubits with either electron shuttling or some electrical interconnect [1609]. In 2026, IQM proposed their Barbell codes, qLDPC variant which accommodates near local couplers with a fixed connectivity [1610]. Another challenge comes with parallelizing logical gates operating on a qLDPC patch [1611].

In June 2026, Jens Eisert et al. broke the consensus on long-range connectivity, arguing that qLDPC could be implemented on planar qubit grids, without any long-range connectivity. They use “directional tile codes”, a family of planar CSS qLDPC codes whose syndrome extraction runs on a square grid using only nearest-neighbour iSWAP/CXSWAP gates. The check-data connectivity is dynamical with nearest-neighbor iSWAP walks both defining the stabilizer supports and implementing their measurement, avoiding the need for long-range connectivity. The space connectivity needed by qLDPC codes is moved into spacetime by iSWAP dynamics during the cycle, so generating longer computing times, meaning another space-time trade-off. However, the proposal presumes that iSWAP is a primary gate in the hardware, and it doesn’t describe how logical operations would be handled.

Quantum Tanner codes are a scalable variation of qLDPC codes with interesting scaling properties [1612–1615].

Bivariate Bicycle Codes which have been selected by IBM in 2025 [1616–1620]. They use weight-6 stabilizer checks, meaning that stabilizer ancilla qubits are connected with two-qubit gates with 6 data qubits, explaining indirectly the relevance of BB codes for the 1-to-6 qubit connectivity planned in the IBM Loon architecture.

Surface codes have been for a long time a reference in QEC, particularly with Google and their 2022 to 2024 experiments on their Sycamore and Willow superconducting processors. They spread single physical qubits quantum information across a large 2D grid of physical qubits to protect it from local noise. It is adapted to planar quantum processors where qubits are connected only with their nearest neighbor, usually 4 neighbors as with Google’s Sycamore and Willow chips, and IBM’s chips since NighHawk in 2025. This square lattice contains data qubits, which hold the quantum state, and measurement qubits, which act as observers, and potentially additional auxiliary qubits.

The whole thing becomes a “logical qubit”. By performing constant parity checks on small clusters of four adjacent data qubits and checking for bit-flips (X-errors) and phase-flips (Z-errors), the surface code helps identify when an error has occurred without actually measuring the data qubits directly. As for each error correction code, these checks, *aka* error syndromes, are used by a classical decoder algorithm to infer a recovery operation.

That recovery identifies an equivalence class of errors modulo stabilizers rather than the unique physical error, and it is usually applied as an update of a classically tracked Pauli frame rather than as physical X or Z gates. The common layout of the surface code is the rotated surface code, which followed the unrotated version as shown in Figure 263.

Encoding the data qubits is not created with CNOT gates like with the textbook examples in Figure 258 and Figure 259, but through parity measurement using X and Z stabilizers (containing a CNOT, CZ or iSWAP), as illustrated in Figure 264 which explains the basics of surface code [1525].

The boundaries of the surface code are used to interconnect it with other logical qubits for the implementation of two-qubit gates using lattice surgery and patches, which are used for routing the connection between distant logical qubits.

Like with every error correction code, the robustness of a surface code is determined by its distance. Here, it corresponds to the size of the physical grid. The distance is the weight of the smallest undetectable logical operator. In the ideal code-capacity picture it bounds the number of errors that can be detected simultaneously ($d-1$) and corrected ($t=\lfloor((d-1)/2)\rfloor$), d being usually odd. Under circuit-level noise, what matters in practice is the logical error rate per cycle achieved by the decoder rather than this counting rule. The $\lfloor$ and $\rfloor$ are floor symbols that rounds a number down to the nearest whole integer. So a distance 3 code can detect two errors and correct one error (flip or phase). A distance 5 code can detect four errors and correct two errors (flip and phase), and so on.

The qubit physical error rate (PER) must stay well below the code threshold, which is around 1% for the surface code under a circuit-level depolarizing model with a matching decoder, and which depends on the noise model, the syndrome extraction circuit and the decoder. With a threshold of 1%, a physical error rate of about 0.1% is needed to obtain a reasonable physical overhead, this factor of ten being an engineering rule of thumb and not a second threshold. Indeed, at threshold, increasing the distance no longer suppresses the logical error rate asymptotically, so no finite code size reaches a given target, which is what the divergence in Figure 265 shows. It explains this, and shows the relationship between the surface code distance and the number of physical qubit per logical qubit. This shows why having one order of magnitude gate fidelities improvement over the threshold is key to obtain a reasonable QEC overhead. Beyond that level, the improvements are more modest as shown in the curve, following a law of diminishing returns.

A QEC round or syndrome extraction cycle is a fundamental unit of time in quantum error correction. A single round is a sequence of physical operations which extracts error syndromes. In a surface code, one round

typically uses four steps: auxiliary qubits initialization, entanglement using CNOT between the replicated data qubits and the auxiliary qubits, X and Z parity measurements with qubit readout, and auxiliary qubits reset. This process is repeated a number of time usually equal to the code distance d . One then looks at the history of the d syndrome measurements to understand where errors occurred during these d rounds, and how to correct the data. This is the role of quantum error syndrome decoders which are classical software running on classical hardware. Some error correction codes with a so-called constant-time overhead do need only one round cycle, at the expense of requiring more ancilla qubits.

qLDPC are competing with surface codes, but these still benefit from continuous improvements, even with cold atoms and trapped ions which benefit from a better connectivity [1630, 1631]. In 2023, Google engineers proposed the yoked surface codes, which divided by two the required number of physical qubits per logical qubits thanks to using small parity checks [1632] and other methods [1633]. Also, in 2026, Google expressed doubts on the high-qubit connectivity needed to implement qLDPC and obtain the space savings advertised in the 2026 Pinnacle paper from Iceberg Quantum [1634–1636]. They estimated that the Pinnacle architecture needed at least 1-to-10 connectivity and that it was not feasible with superconducting qubits. IBM thought as well that their Loon architecture with a 1-to-6 connectivity was already a challenge. Google then improved their yoked codes technique in May 2026 with their “denser planar surface code” which further reduces the physical overhead of surface codes by a factor 4.5 [1637]. In September, though, Iceberg Quantum, published a new resource estimate with a 1-to-8 connectivity, landing in the same zone with a month of computing time with 120,000 physical qubits [1638].

In China, researchers from Hefei, Hangzhou and Guangzhou prototyped in 2026 on 125 superconducting qubits two distance-7 logical qubits arranged on the same patch, to enable the transversal state teleportation between both logical qubits [1639]. This is however an exception. Another team in China, led by Jian-Wei Pan, proposed in 2026 its “hierarchical logical processor” (HLP) approach which combines surface codes and a new type of transversal CNOT gate, that they call the “hybrid-unit CNOT gate” [1640]. It couples an elongated RSC (“rotated surface code”) patch (*aka* a shuttle bus) to multiple RSC patches, or cores, simultaneously. They found this method to be more efficient than the yokes codes and planar surface code from Google.

XZZX surface code that would reduce the number of required physical qubits to create a logical qubit thanks to a lower error threshold. In April 2021, University of Sydney science undergraduate Pablo Bonilla Ataides XZZX paper published in Nature Communications brought the attention of Amazon researchers [1641]. This surface code could be used by Amazon who made a choice to use a relatively low number of




Figure 263: Unrotated vs rotated surface code. The rotated surface code needs about 75% of the physical qubits used in the unrotated surface code to create a logical qubit with the same error-correcting distance. Source: [1621].




Figure 264: Distance 3 rotated surface code with 9 data qubits (white) and 8 auxiliary measurement qubits (black) [1621]. CNOT gates handle local interactions between the data and auxiliary qubits. The auxiliary qubits in the blue zones detect X error in the Z stabilizers connected to 4 data qubits, and the auxiliary qubits in the red zone detect Z errors in the X stabilizers also connected to 4 data qubits. The Hadamard gates in the X stabilizers are used to interchange the Z and X basis. There are weight-four (inside the rectangle) and weight two stabilizers (on the boundaries of the code). The plaquette and star naming, like the colors used here, is a convention attached to this particular layout and syndrome extraction schedule. Color alone does not establish the stabilizer type, which follows from the stabilizer operators themselves, the ancilla preparation and measurement bases and the order of the entangling gates, all of which have to be read off the circuit gate by gate. It contains d^2 data qubits + $d^2 - 1$ auxiliary qubits = $2d^2 - 1$ total physical qubits. The “bootstrapping” sequence in a FTQC quantum computer resets all physical qubits to $|0\rangle$. Then the logical qubits are initialized by running the stabilizers which will entangle the data qubits from each logical qubit. Then, logical gates are implemented, and between these, cycles of error correction are implemented in each logical qubit. In the common case where the primary gate set contains only a CZ and not a CNOT, an H gate is applied beforehand on the auxiliary qubit for the Z stabilizers, and applied before and after the entangling gates on the data qubits for the X stabilizers. When the primary two-qubit gate of a quantum processor is an iSWAP, QEC are designed to be iSWAP-native and they require the same number of iSWAP gate that their equivalent using a CNOT or CZ gate [1622, 1623]. A variant of the surface code is the XZZX surface code where all tile stabilizers are the same, measuring X and Z errors with a single auxiliary qubit, with an HXX gate applied for the Z error detection for 2 of the 4 data qubits. This variant is well adapted to biased noise qubits where phase noise is more important than flip noise. Also, error syndrome detection is much easier [1624]. Source: [1625].

photons per cat qubit (8 to 10, compared to about 20 for Alice & Bob [1642]), still requiring some first level bit-flip error correction on top of phase-flip correction. That’s where an XZZX surface code QEC could come into play. These codes are indeed mentioned as an option in Amazon’s technical paper from December 2020. A team from Amazon, Caltech and the University of Chicago improved these codes in 2022 to work with biased noise (which has more phase than flip noise) [1643, 1644]. The established name is the XZZX surface code,

and the ZXXZ labeling found in some papers designates the same family up to a relabeling of the bases, which is a local Clifford transformation combined with a choice of boundary and noise conventions, and not simply a 90° geometric rotation of the lattice. It was used by Google in its 2024 Willow break-even experiment, with logical qubit gates having longer coherence time than physical qubits [1645]. This code is adapted to biased noise physical qubits which exhibit more phase than flip noise, with $\eta = p_z/(p_x + p_y)$ (often approximated

Q Differentiating surface codes and qLDPC

One key difference between surface codes and qLDPC codes is that a surface code encodes a single logical qubit in a code patch. It's a $[[n, 1, d]]$ code. qLDPC encode multiple logical qubits in a patch and data qubits are shared between multiple logical qubits. For example, a $[[144, 12, 12]]$ qLDPC code encodes 12 logical qubits with 144 data qubits, here, with a distance 12.

Surface codes are a subset of qLDPC codes with $k = 1$. “Low-Density Parity-Check” in qLDPC means that each parity measurement only checks a small, fixed number of qubits which, in a surface code, is usually 4. Quantum information researchers use “qLDPC” to label “high-rate” qLDPC codes. Shifting from surface codes to qLDPC means abandoning the 2D nearest-neighbor topological constraints where $k = 1$ in favor of long-range, non-local qubit connectivity that enables $k > 1$.

Some qLDPC families offer a reduced time overhead, the surface code logical cycle of d measurement rounds being replaced by fewer rounds, down to a single round for single-shot constructions. This

is a property of particular families and noise models rather than of qLDPC codes in general, since measurement faults usually require either repeated rounds or a single-shot construction with its own assumptions, check weights [1626] and additional ancillas. This shift has some trade-offs as you obtain a better encoding rate and need fewer physical qubits, and, for some families, fewer syndrome rounds, at the expense of greater connectivity requirements between the qubits and, usually, of more expensive classical error syndrome decoding.

qLDPC are popular in two main cases: with atomic qubits which natively enable long-range qubit connectivity, usually with atoms/ions shuttling, or with planar solid-state qubits where an additional connectivity layer is added in the chip to enable long range connectivity, and which can be optimized to minimize the number of connectivity layers as described by the MIT in 2026 [1607]. Even silicon qubits can accommodate electron shuttling to enable this long range connectivity, and in some cases, with shuttling parity check qubits instead of data qubits [1627].

to p_z/p_x) sitting between 10 and 100 like with some transmon qubits, with p_z being the phase-flip probability and p_x the bit-flip probability). η can exceed 10^5 for cat-qubits which have a marginal flip noise and rely mainly on 1D repetition codes, at least for supporting circuits in the megaquop range.

Color codes are topological codes that enable transversal gates working with a single step across the entire qubit lattice, including H and S gates which are hard to implement directly with surface codes. They require fewer physical qubits due to a need for d^2 qubits for a distance d and they are faster to operate. They are more sensitive to physical qubit noise than the surface code, having a lower threshold, but their error syndrome decoding cost is higher [1646–1649]. Among others, Google is working on it [1650] as well as RIKEN [1651].

There are many variations of color codes with 2D and 3D codes. Their plaquette shape can be triangular (on the edges) or hexagonal, or even octagonal (inside it). Color codes are laid out on a trivalent graph with 3-colorable faces (red, green blue) with all adjacent shapes being of a difference color. They generally use weight- n stabilizers, n depending on the face shape, with every face supporting both an X and a Z stabilizer. Figure 266 shows the basics of color code operations with the related circuit for a simple distance-3 triangular color code with four zones. There is also a “domain wall color code” variant [1652]. The 50% figure sometimes quoted for it cannot be a circuit-level threshold: such a value only makes sense for a very specific code-capacity or

erasure-type noise model, and the noise model, decoder and boundary conditions have to be stated alongside it.

Figure 267 shows a more detailed representation that includes auxiliary qubits on the right [1655].

Figure 268 shows a maze of lattice surgery patches to link color codes together.

Color codes are envisioned for multiple qubit platforms, from superconducting qubits to trapped ion and cold atom qubits.

Quantum rainbow codes is an extension of color codes and pin codes which enables transversal non-Clifford gates (T) [1657].

Floquet Codes *aka* Planar Honeycomb Codes, were introduced by Matthew B. Hastings and Jeongwan Haah from Microsoft in 2021 to simplify toric codes with fewer qubits and stabilizers [1658]. These codes are adapted to qubits architectures which do not have a native CNOT gate and instead use joint pair-wise qubit measurements on two qubits (XX, YY, ZZ) like with MZMs at Microsoft and with PsiQuantum measurement based computation. Using surface codes with these qubits shows lower performance because the syndrome extraction circuit is more complicated [1659, 1660]. The technique was improved by Google researchers [1661], by Christophe Vuillot from Inria Nancy (now at Alice & Bob) [1662, 1663], by Riverlane [1664], with the XYZ² variant in 2022 [1354], by the MIT and Microsoft with so-called dynamic automorphism codes [1665], by IBM




Figure 265: How the number of physical qubits per surface code logical memory qubit relates to the surface code distance and with the ratio between the physical error rate (PER) vs the threshold rate. The curve diverges as the physical error rate approaches the threshold: at threshold, a larger distance no longer suppresses the logical error rate, so no finite code size reaches a given target. You get the best gain with an error rate about ten times below threshold, which is an engineering rule of thumb and not a second threshold. Afterward, there is a law of diminishing returns when the gain in fidelities decreases much slower the number of required physical qubits [1628]. p and p_L need to be carefully described and are not really equivalent. These correspond to physical and logical error per cycle of operation. One full cycle of surface code syndrome extraction is broken down into a sequence of discrete time steps during which every qubit is doing exactly one of five things: being initialized or reset (state preparation), running a single-qubit or a two-qubit gate, being measured (readout) or doing nothing (idle operation). p corresponds to the “single-step error rate” assuming that all five of these operations take the exact same amount of time and have the exact same error probability, which is a simplification. p_L is a logical error rate for the logical qubit memory, regardless of the logical operations executed on this memory. It is an error rate for the logical qubit per error correction cycle round. It doesn't correspond to a logical operation error rate. It corresponds more about the capability to create virtual logical qubits with a longer coherence time than the underlying physical qubits. Sources: [1629], and compilation by Olivier Ezratty, 2021-2026.

[1666], by AWS [1667], even NuQuantum [1668], and in China [1669].

Quantum Reed–Muller (QRM) codes are CSS codes built from classical Reed–Muller codes $RM(r,m)$, with X-type stabilizers from $RM(r,m)$ and Z-type from the dual $RM(r-1,m)$ (Steane, [1670]). Their special feature is the support of a transversal non-Clifford gate, with the famous transversal T gate of the $[[15, 1, 3]]$ code, which explains its use in magic state factories and code switching. The small members are the $[[7, 1, 3]]$ Steane code and the distance-3 family $[[2^{m-1}, 1, 3]]$. There of course many variations of such codes, like [1671].

Other error corrections code variations are:

Continuous error correction. In opposition to discrete quantum error correction codes (DQEC) that we've covered, there are alternative QEC using continuous measurement and correction. Juan Pablo Paz and Wojciech Zurek proposed in 1998 a continuously operating error correction code, the CTQEC, for “continuous-time QEC”, or CQEC, based on differential equations and acting at reduced time intervals. There are two

methods for acting directly on the information (direct CTQEC) or via auxiliary qubits (indirect CTQEC). In the direct variants, stabilizer information is obtained by weakly measuring the physical qubits rather than through projective ancilla readout, while the indirect variants still use probes, ancillas or engineered reservoirs. Claims of faster detection or of deferred recovery are protocol-dependent and rest on stated assumptions about the measurement, the estimator, the feedback bandwidth and the noise model. These methods were later improved by various contributors including Andrew J. Landahl and Gerard J. Milburn [1672–1674]. The latter proposed to use some real-time measurement-based estimator (MBE) of the real logical qubit to be protected to accurately track the actual errors occurring within the real qubits in real-time. This leads to the **MBE-CQEC** scheme that protects the logical qubit to a high degree and allows the error correction to be applied either immediately or later.

Dynamical codes, in current QEC usage, codes are defined by a time-dependent schedule of measurements

Q Pseudo-threshold and break-even

These two terms describe related but distinct milestones, and a third one, below-threshold scaling, is often mixed up with them. This is the easiest target because the bar is relatively less demanding.

A **pseudo-threshold** is the theoretical physical error rate (PER) where a small, specific logical block, with a given code distance (like $d = 3$) first achieves a lower logical error rate per round than a physical qubit executing the same idle or gate operation [1653]. But a pseudo-threshold varies and usually increases with the code distance.

The **break-even level** corresponds to the experimental realization of the pseudo-threshold. It is when demonstrating that the QEC overhead is finally doing more good than harm for a specific finite-sized logical qubit.

Below threshold is another criterion which means that the logical error rate per cycle decreases as the code distance grows, under a given noise model. With Willow in 2024, Google reported below-threshold surface code memories, the distance-7 memory showing a lower logical error rate per cycle than the distance-5 one [1645]. Any break-even claim must in addition name the physical baseline it is compared against.

We've seen in Figure 265 that the threshold was not a sufficient level to create a workable logical memory. Practically speaking, you need to reach an average physical error rate (PER) at about 1/10 of the threshold to implement practical error correction with a reasonable physical overhead for large code distances. This is however just an engineering rule of thumb.




Figure 266: A visual representation of a distance-3 color code (left), the logical qubit initialization (middle), and the stabilizers circuit with X and Z correction (right). Each zone in this simple version contains four data qubits and internal auxiliary qubits, or a combination of an auxiliary qubit and a flag qubit (usually not shown) which handle the X and Z stabilizers. Color codes can use more than one auxiliary qubit per zone. Source: [1654].

or gauge choices rather than by a fixed set of stabilizers. Their logical operators are spacetime objects. Floquet codes are the best known example. They are not defined by their reaction to time-varying noise.

This is a different topic from adaptive calibration and noise-aware decoding, where the correction strategy is adjusted to noise that evolves over time and exhibits temporal correlations, instead of assuming errors that are independent and identically distributed. That second family uses feedback mechanisms where syndrome measurement results influence subsequent operations, real-time noise characterization, and decoders whose priors are updated from measured noise. It draws on quantum control theory and non-Markovian noise models, which are well studied for qubit types like cold atoms [1675], and it requires fast and accurate noise characterization, real-time processing capabilities, and efficient

feedback mechanisms [1676, 1677]. Both notions can be applied to surface codes, qLDPC codes and other codes.

Another approach uses machine learning techniques to predict the evolving noise patterns and adjust the error correction strategy accordingly during computing and not only in post-processing tasks as is done with quantum error mitigation (QEM).

“Whatever comes out of these gates, we have a better chance to survive if we work together. You understand? We stay together, we survive.”

General Maximus Decimus Meridius (Russell Crowe) in *Gladiator*, 2000.




Figure 267: A generic color code of distance-5 with its colored cells (left) and variants of a fault-tolerant distance-5 and distance-7 color code using 3 auxiliary qubits and 3 flag qubits inside the colored cells (right), which implement a X and Z weight-6 stabilizer. Source: [1655].




Figure 268: Lattice surgery example with color codes. It shows that lattice surgery can have a significant overhead in physical qubits. FTQC resources estimates must account for this, on top of the cost of magic states creation and teleportation. Source: [1656], which, compared to surface code lattice surgery, generates a 3× improvement in the space-time overhead, with a combination of a 1.5× improvement in spatial overhead with a 2× speedup thanks to the parallelization of commuting logical measurements.

Spacetime quantum error correction codes incorporate both spatial and temporal dimensions to protect quantum information more comprehensively, by correcting errors occurring over time and space [1574]. It is implemented in MBQC QEC.

Autonomous QEC implements error correction at the hardware level with various forms of engineered dissipation [1678]. The most commonplace aQEC are **bosonic codes**, including GKP error codes [1679, 1680], binomial codes, 0- π [1681], cat-codes or cat-qubits. Cat-codes are used by the startup Alice&Bob. Knowing that their creation goes back to the work of Mazhar Mirrahimi and Zaki Leghtas in 2013, with whom the founders of Alice & Bob worked. Error correction codes are constantly being updated. A proposal from 2019 goes further than cat-code and does not depend on hardware architecture [1682] and spherical codes which extend cat-codes [1683]. The latter implements in a cavity two super-

posed Schrödinger cats that manages a projection space used for error correction. It usually corrects flip errors autonomously with T_1 reaching 10 seconds [1684]. New variations can correct both flip and phase errors [1685]. It must also be engineered to correct entanglement errors [1686]. Qutrits can also be built out of bosonic qubits, extending the computational space [1687, 1688]. Bosonic qubits can also rely on using two transmons to actively correct single-photon loss and low-frequency dephasing [1689].

Hardware specific QEC codes also exist like for spin qubits which have significant phase errors [1690, 1691], photonic qubit which rely on measurement based quantum computing and entirely different QEC schemes [1692, 1693], topological qubits which are well protected against errors but not entirely and still require some QEC [1694], neutral atoms [1695], trapped ions [1696,

[1697], Majorana fermions [1698], and also superconducting qubits [1623, 1633, 1699].

2D, 3D and 4D codes. The most commonplace quantum error correction codes use 1D or 2D physical qubit layouts. Surface codes with superconducting qubits are 2D codes using nearest neighbor qubit connectivity, although they have 3D and 4D variants [1700]. qLDPC codes use long range qubit connectivity and can leverage many-to-many qubit connectivity brought by various qubit modalities like trapped ions, cold atoms (with atoms shuttling) and even superconducting qubits (with long range connectivity like with the IBM C-couplers used within qubit chips). QEC codes are defined not just by their algebraic structure, but by the geometry and topology of the complex that defines their checks. That dimension labels the code's defining cell complex or locality structure and not necessarily the physical arrangement of the qubits, since a code can be implemented in another geometry at the price of routing, non-local couplers or spacetime overhead.

- **2D codes** arrange qubits on a 2D grid and uses local measurements to detect and correct errors. The most famous such code is the surface code [1549].
- **3D codes** arrange qubits on a three-dimensional lattice. Error correction involves detecting defects such as loops or surfaces in this 3D space. These codes can sometimes be self-correcting, meaning they naturally resist errors without constant intervention [1701]. qLDPC codes are defined by sparse parity check matrices and many-to-many connectivity between qubits and stabilizers. While these codes can be implemented in 3D physical architectures, they are not inherently 3D topological codes unless their error correction mechanism relies on 3D spatial topology [1702]. Examples of such 3D code are the Layer Code [1703] and the 3D STC code [1704].
- **4D codes** use a four-dimensional lattice, often modeled as a tesseract or hypercube [1705]. The fourth dimension refers to the code's defining complex and not to a physical arrangement of qubits. Implementing such a code requires either non-local connectivity or additional routing and time overhead in a 2D or 3D layout. These codes allow for single-shot error correction and support loop-only toric codes and Hermite-normal-form lattice rotations [1706]. Microsoft's 4D quantum error-correcting codes presented in 2025 are built on a 4D torus lattice [1707, 1708]. This structure is periodic in all four dimensions and supports loop-only toric codes. Logical operators correspond to closed loops in 4D space.

Currently, no qubit modality has demonstrated a fully functioning 3D qubit lattice with local control and error correction. Current superconducting qubits use flip-chip bonding for control routing, not for 3D qubit layout. Neutral atom arrays can trap atoms in 3D configurations,

but full local control in 3D is not yet realized due to limitations in tweezers and gate operations. Trapped ions can be manipulated in 3D traps, but scalable 3D lattices with local gates are not yet feasible. Finally, photonic qubits are typically confined to 1D or 2D waveguide structures [1709]. Whatever, interactions respecting the logical structure of a 3D or 4D lattice can, in principle, be realized in 2D physical layouts if sufficient connectivity between the physical qubits is possible.




Figure 269: Microsoft representation of its tesseract 4D error correction code. Source: [1707].

Here are now several QEC features to deal with:

Leakage errors corrections. Most quantum error correction codes are focused on Pauli channel noise errors correction (flip, phase and depolarizing errors). Leakage errors require their own correction methods, which require additional qubits, although in reasonable numbers like the 4 leakage unit qubit in Google's 105-qubit Willow chip [1710].

Erasure errors are characterized operationally by the fact that the decoder knows which qubit was affected. Leakage and loss, where the physical object lands outside the computational subspace or disappears, become erasures only once a detection mechanism reliably flags and locates the event without revealing the encoded logical information. Without that flag, they remain ordinary uncorrected errors with their own dynamics. It can be an atom loss with trapped ions and neutral atoms or a shift of a transmon qubit to a higher energy level than the ones corresponding to states $|0\rangle$ and $|1\rangle$. Leakage Reduction Units (LRU), Leakage Detection Units or Loss Detection Units are used to convert these errors into erasure errors that are detected without damaging functional qubits, particularly with cold atoms [1711]. Other solutions exist using an erasure decoder [1712]. Atom losses can also be detected using graph neural networks applied to error syndrome detection data [1713].

Erasure conversion is about tuning a quantum processing unit so that dominant errors are erasure errors, with so-called erasure qubits. Erasure errors occur when the qubit exits its computational space in a state that can easily be detected without damaging qubits which

are still in a superposed state from the computational space. For example, cold atoms can decay to a state that does not correspond to the states $|0\rangle$ and $|1\rangle$ and which can easily be detected [1714–1717]. This is also used by Quantum Circuit Inc with its dual-rail superconducting qubits (the company was acquired by D-Wave in January 2026). The technique can then be integrated with qLDPC codes, with long-range atom interactions using multiple lasers to enable CZ gates between distant qubits [1718, 1719]. Erasure error conversion can also be implemented with dual rail superconducting qubits using pairs of transmons where the states $|0\rangle$ and $|1\rangle$ correspond to having a microwave photon in one of the transmon cavities. In that case, the dominant error is a loss of this photon which can easily be detected without damaging a coherent superposition of 01 and 10 [1720–1723]. Detecting erasure errors is a cheaper alternative to complicated error syndrome decoding on large surface codes or equivalent codes. This error conversion can reduce the threshold error of the qubit and error correction code.

Flag qubits are extra auxiliary qubit used in fault-tolerant circuits to signal the occurrence of specific error patterns that could lead to high-weight, uncorrectable errors during stabilizer measurements [1724–1726].

Qubits lifetime extension. If we need to accumulate error correction codes, don't we risk running into the wall of qubit decoherence, particularly with superconducting qubits? Well, no. As said before, error correction codes have the direct effect of artificially extending the coherence time of the qubit registers by several orders of magnitude [1727]. Repeated syndrome extraction and decoding suppress the logical error rate below threshold and extend the lifetime of the encoded information. They do not reset the physical T_1 and T_2 of the underlying qubits, which keep decohering at their material rate throughout the computation. This is why the 2019 Google resource estimate for Shor's integer factoring, with 20 million physical qubits and 8 hours of run-time, can exceed the physical coherence time by many orders of magnitude. Under the assumptions of that estimate, the information is continuously encoded, diagnosed and refreshed at the logical level, not because a physical coherence clock is reset. A record qubit lifetime doubling extension to 1.8 ms was achieved by a Yale and Sherbrooke University team led by Michel Devoret in 2022 with superconducting qubits [1728]. It was a GKP bosonic mode coupled to a tantalum transmon and using reinforcement learning to implement an autonomous error correction. It was not really a logical qubit made of several identical physical qubits.

Concatenated codes is about exploiting the recursivity of some error correction codes like qLDPC [1729]. A QEC code generates logical qubits which can then be used as virtual physical qubits for a new QEC, and so on. With each recursion, the apparent error rate decreases (Figure 270). We stop concatenating QEC codes when we reach an error rate compatible with the

circuit size. Per the famous threshold theorem, concatenation can be optimized by using different types of QEC at each level of recursivity [1730, 1731]. This theorem was demonstrated only for a 7-qubit Steane error correction code and for error rates that are not growing with the number of physical qubits. This is unfortunately not what is currently observed with most qubit types! Surface codes and their various derivatives are not concatenated but rather expanded in 2D with a growing number of qubits. But their relative noise and number of qubit scaling are different.

With concatenated codes built from a distance-3 block, the logical error rate follows a recursion of the form $(\epsilon/\epsilon_{th})^{2^k}$ with a number of qubits in X^k , ϵ being the physical error rate, ϵ_{th} the threshold of the gadget used, k the number of concatenation levels and X the number of qubits within a single concatenation. This doubly exponential form depends on the code, the fault-tolerant gadgets, the decoder and the noise model and is not a universal law [1733]. The number of physical qubits depends on the implementation and on the way auxiliary qubits are managed and recycled, using Steane's method (not to be confused with Steane's code). In [1734], 91 qubits are used with a Steane 7-qubit $[[7, 1, 3]]$ code. It includes the auxiliary factory and $4 \times 7 = 28$ qubits auxiliary factory times 3 because the preparation, verification and measurement of auxiliary qubits is three times longer than the data qubit operations (9 versus 3 time-steps). Hence, while auxiliary qubits are used for a given gate, they must be prepared for the next two gates. Thus each level of error correction replaces one qubit by 91 qubits (7 data qubits and 3×28 auxiliary qubits).

Concatenation of codes C_1 (size n_1) and C_2 (size n_2)

We construct a code of size $n_1 n_2$, where each qubit of C_2 is replaced by a block of n_1 qubits encoded in C_1 .

Higher order QEC by concatenation

Level of concatenation	Error probability
Physical qubits	$\epsilon_0 = p$
1 st encoded level	$\epsilon_1 = c p^2 = c^{-1} (c p)^2$ (*)
2 nd encoded level	$\epsilon_2 = c (c p^2)^2 = c^{-1} (c p)^{2^2}$
⋮	⋮
r th encoded level	$\epsilon_r = c (\epsilon_{r-1})^2 = c^{-1} (c p)^{2^r}$

(*) For the Steane code $c \approx 10^4$

Figure 271: How concatenated codes are reducing the logical error rate depending on the number of nested concatenations. Source: [1397].

Flag qubits can lighten this significant overhead and reduce X . This doesn't take into account T gate magic state distillation. The number 15 refers to the 15-to-1 protocol, which consumes 15 noisy logical magic states per accepted output, and not to a universal minimum of 15 physical qubits, the physical overhead depending on the encoded blocks, the ancillas, the routing and the repeated checks [1735–1739]!

Concatenated quantum computation:




Figure 270: How QEC concatenation works with level-0, level-1 and level-2 CNOT gates. Source: [1732].

Thresholds mentioned for concatenated codes vary by several orders of magnitude depending on the code, the gadgets, the locality assumptions and the noise model, the most conservative rigorous proofs landing around 10^{-6} [1733] while later constructions report much higher values. These numbers are properties of a specific construction and not of the concatenated class. With surface codes, the noise is reduced according to $\epsilon^{d/2}$ and the number of required physical qubits grows by d^2 , d being the distance of the surface code, more or less corresponding to the edge of the surface code squares for distances 3 and 5.

In summary, concatenated error correction codes have a better impact on noise, but at the expense of a large number of physical qubits while surface codes scale slower in error reduction and physical qubits requirements [1538]. It seems that surface codes are more appropriate for more noisy physical qubits while concatenated codes will be better, for less noisy qubits. Thanks to surface codes, the thresholds quoted in the literature have moved from the 10^{-6} of the earliest concatenated constructions to about 10^{-2} for a surface code under a circuit-level depolarizing model with a matching decoder. These two numbers come from different codes, gadgets, noise models and decoders and are examples rather than class constants.

Heterogeneous Quantum Error Correction (HQEC) uses multiple codes that are applied simultaneously or sequentially to counteract various error types without needing to switch, and not dynamically like with code switching [1740].

Noise models are important tools to select the right quantum error correction codes on a given system. They modelize all error types on a qubit processor and their probabilities. They define the various error channels impacting the qubits (see a simplified example error model in Figure 3 in [1741]). They specify for a gate, a reset, an idle step, or a measurement, which Pauli errors can occur and with what probability. One such error model is well documented in the Supplemental Materials of [1135], from QuEra (page 3) with, for example, 4×4 error channel matrices for entangling gates. They can help create a digital twin that simulates the effect of error correction and fault tolerant techniques.

Noise models can be completed by detector error models (DEM) which are adapted to error correction decoders. They are maps that create a correspondence between the detected syndrome errors and the errors themselves. A DEM is created using the physical qubits noise model [1742]. A DEM can be created manually or automatically, as proposed in [1743] with LightStim.

Fidelities space and time variability. As industry vendors are publishing ambitious plans for supporting large circuits in the gigaqubit to teraqubit range, some cautiousness should be used for at least two reasons.

One is that the noise models used to create QEC/FTQC resource estimates are often based on using a simplistic fidelity levels. In most cases, the basic assumption is that the FTQC hardware supports 99.9% fidelities. It means that all operations fidelities (reset, single and two qubit gates, readout) have a stable fidelity of 99.9%. Real world systems do not behave like this. Error rates benchmarks show that errors follow a distribution centered around the advertised values, and with some standard deviation. The provided values are median or average values and they can fluctuate in space and time. The impact of large standard deviation around median values on the cost of error correction must be understood [1744]. Another limitation of current QEC/FTQC models is that they do not take into account that quantum gates and readout operations used in QEC are themselves generating errors [1745].

The good news is that in 2026, [1746] demonstrated that there are some boundaries of acceptable defectiveness in qubit fidelities, but that sufficiently large code distances can cope with the variation. [1747] provided in 2017 some numbers on the acceptable rate of defective qubits on a chip and computed effective code distance, meaning that it grows with the qubit defect rate. In its 2024 Willow surface code experiment [1645], Google did show that chip non-uniformity was not catastrophic when the system fares below QEC threshold. And likewise, [1748] did show that lattice surgery don't have the same constraints with two-qubit gate fidelities, opening the realm of distributed quantum computing with less stringent constraints on Bell pair distillation.

So, it is also important to characterize the variability of these fidelities over time and space in the quantum

processor [1352, 1749, 1750], as shown in Figure 272. Calibration is used to detect faulty qubits and faulty two-qubit connections so that they are avoided by the compiler [1751]. Also, sometimes, industry vendors mention that their two-qubit fidelities can reach 99.9%, which is misleading since it doesn't represent a median or standard deviation which are the two things that count for implement some viable QEC/FTQC approach. In that case, the median value is never 99.9% but much lower. And by the way, as of September 2026, we are not aware of a commercial system reporting a median assignment fidelity above 99.9% under simultaneous readout of a full register of more than a hundred qubits. Selected-qubit or SPAM figures above that level are reported on small registers, which is a different claim.




Figure 272: The difference between the fidelities assumptions in most QEC/FTQC industry vendor resource estimates, particularly for breaking RSA-2048 keys, and their actual distribution. The distributions drawn here are conceptual and are not plotted from a real dataset, and reset, single-qubit, two-qubit, readout, leakage, idle and correlated errors each have their own distribution, which also depend on the qubit type and implementation. There is some uncertainty on the evolution of these distributions as the system size grow. Current readout fidelities are usually lower than single and even two-qubit gate fidelities. On one hand, they will benefit from various scientific engineering improvements which may increase fidelities, but on the other hand, as system size grows, new error sources may counter that effect. Only experiments will tell! © Olivier Ezratty, September 2026.

Second is that most QEC/FTQC schemes are focused on correcting simple errors like Pauli stochastic errors while errors on real-world qubits and control operations are not Pauli stochastic. These errors can change QEC/FTQC thresholds, increasing the space-time overhead of QEC and magic state factories. It can also increase logical error rates [1752]. As a result, we know that, as FTQC systems will progressively scale, error models will need to be refined as well as QEC/FTQC overhead estimates. And this considers that correlated errors are well handled.

Circuit level noise model corresponds to a noise model that injects errors at every physical operation in

the actual stabilizer measurement circuit including gates, measurements, idle states, and state preparations. It is the most realistic and demanding noise model used in QEC digital simulations. It can help generate detector error models (DEMs).

Diamond norm, *aka* diamond distance, is a mathematical model used to quantify the difference between an ideal quantum operation and its noisy, physical implementation [1181, 1744, 1753]. It provides a worst-case bound on how distinguishable a noisy channel is from an ideal channel. This norm is applicable to the execution of a whole quantum circuit. The diamond norm possesses two properties: it is stable under entanglement and covers the worst-case scenario allowing the system to be entangled with an arbitrary ancillary state, and it exhibits subadditivity under composition. This means that when chaining gates, the diamond norm guarantees that the total error of the circuit grows linearly at worst. The error of applying the sequence of gates is bounded by the sum of their individual diamond distances from the ideal gates. In that context, a diamond-norm error budget is an architectural and algorithmic planning tool which takes the maximum allowable failure probability for an entire quantum circuit and mathematically distributes that allowance across all the individual logical operations required to execute it. Compared to a circuit noise model, the diamond norm model is about evaluating the worst case scenario while the circuit noise model is an operational description of what errors happen and where they occur in a circuit.

To summarize the QEM, QED and QEC above text, Table 8 contains a summary of the key points differentiating these, given that some quantum error mitigation techniques may be used in the FTQC regime along with quantum error correction codes.

5.2.8 Error syndrome decoding

Error correction codes like surface codes, color codes, and qLDPC's primary goal is to detect error syndromes using data extracted with parity measurements. These parity measurement results must be computed to detect actual errors, where they occurred and how to correct them with single quantum gates [1756]. This is the task allocated to "error syndrome decoding", a classical computing task that must be run real-time for at least two reasons: errors must be corrected as fast as possible otherwise the correction may be wrong due to the dynamics of the system ("drifting noise" [1757]) and it must also cope with the speed of the logical gates executed on the QPU.

Finding a maximum-likelihood correction is NP-hard in general, but practical decoders do not solve that problem and run in polynomial or near-linear empirical time. What matters operationally is decoder throughput, latency, the feed-forward deadline and the backlog, rather than worst-case hardness. Many techniques are being devised both at the software and hardware level to op-

	QES quantum error suppression	QEM quantum error mitigation	QED quantum error detection	QEC quantum error correction
goals	prevent errors at the pulse and gate level, before they accumulate	reduce bias in estimated expectation values through classical post-processing	detect errors with distance-2 codes and discard corrupted runs with post-selection	real time protection of logical qubits, supports large circuits
used techniques not exhaustive	dynamical decoupling, shaped pulses and optimal control, Pauli twirling [1754]	ZNE, PEC, Clifford data regression, symmetry verification [1413, 1418]	error-detecting codes such as distance 2 $[[4, 2, 2]]$, flag qubits, stabilizer and parity checks [1755]	surface codes, color codes, qLDPC codes, Floquet codes, Steane and CSS codes, real-time decoding [1549]
main figures of merit	coherence time (T_2) gain, gate fidelity improvement, no qubit overhead	sampling (shot) overhead, classical post-processing cost, bias reduction	acceptance and discard rates, post-selection overhead, output fidelity gain	logical error rate, physical-to-logical qubit ratio, code distance, decoder latency
space and time overheads	extra control pulses and calibration, no ancilla qubits, no post-processing	more circuit shots which grows exponentially with circuit volume, plus classical post-processing	discarded shots with acceptance dropping exponentially with size, ancillas qubits and mid-circuit measurement	physical qubits per logical qubit overhead, fast real-time decoding, longer runtime
where	NISQ and FTQC	NISQ	early FTQC	FTQC

Table 8: Handling quantum errors: suppression, mitigation, detection, and correction. These four families are complementary rather than a strict ladder, they can be combined, and the overheads quoted in the table belong to the specific protocols named in each cell rather than to the family as a whole. They are ordered here from left to right by increasing capability, overhead, and architecture or time horizon from left to right. © Olivier Ezratty, September 2026, updating a figure from UQT 2025 which compared only QEM and QEC.

timize syndrome detection and avoid this curse [1527]. Syndrome decoding is mostly about solving complex graph problems since the qubit to correct depend on the layout of the error syndrome on the (usually) 2D logical qubit’s physical qubit layout.

Error syndrome decoding usually makes use of a Tanner graph which is a graphical representation of the structure of a code, to facilitate the creation of decoding algorithms like belief propagation. It is a bipartite graph consisting of variable nodes representing physical qubits and check nodes corresponding to parity-check stabilizers. The Tanner graph’s edges connect variable nodes to check nodes if the corresponding qubit participates in that parity check in the error correction code (see examples of Tanner graphs in [1630]).

Here are the main types of error syndrome decoders:

- **MWPM** (minimum-weight perfect matching or minimum-weight decoders) is an algorithm used in surface codes that interprets the syndromes and determines the most likely error configuration that was detected [1758]. It involves finding the minimum-

weight matching of the syndrome graph, which is a bipartite graph representing the connections between the syndrome qubits and the corresponding data qubits. The graph’s edges are assigned weights based on the syndromes obtained from the syndrome measurement. The weights indicate the likelihood of an error occurring between the connected qubits. The algorithm finds a matching that minimizes the total weight of the edges, which corresponds to the most probable error configuration. Once obtained, the identified errors can be corrected using recovery single Pauli gate operations. Minimum-weight perfect matching is itself solvable in polynomial time, and implementations report empirical scalings from $O(d^3)$ to $O(d^5)$ with a surface code of distance d , depending on the variant and the noise model. It is exact maximum-likelihood decoding, a different problem, that is hard in general, and, of course, researchers are trying to minimize this significant classical computing overhead [1759–1761]. Scaling can be further improved thanks to parallelization [1762] and with the help of graph neural networks [1763]. There is a

MWPM variant adapted to bivariate bicycle codes [1764].

- **Most Likely Error** (MLE) decoding which relies on approximation and heuristics, works with qLDPC, but scales badly. But variants like Minimum-Weight Parity Factor (MWPF) are more efficient [1765]. Google developed its own variant of MLE, the Tesseract decoder which is applicable to multiple QEC codes (surface codes, color codes, qLDPC) [1766, 1767].
- **Union-find decode** (UF) is a syndrome decoding technique that doesn't find the perfect matching of error configuration but is more efficient and faster than MWPM since it scales as $O(d^2)$ [1768] to $O(d^3)$. Its running time scales linearly with the code size and can work on surface codes and qLDPC [1769]. It is adapted for real-time decoding and FPGA implementations. A standard Union-Find decoder guarantees correction of errors up to half the code distance, while a greedy decoder finds errors immediately based on local proximity, which is faster but doesn't work well on worst-case scenarios. In 2026, [1770] provided the first rigorous proof that both standard and greedy UF decoders can achieve a finite fault-tolerance threshold under realistic circuit-level noise, proving that they can scale on hardware despite the greedy decoder lacking adversarial guarantees.
- **Belief Propagation decoders** is a class of syndrome decoders that are efficient for decoding qLDPC codes due to their iterative message-passing nature [1771–1773], although some versions are also adapted to surface codes decoding [1774]. They can be parallelized and combined with other techniques like Union-Find and MWPM decoders, or machine learning techniques [1775]. But their convergence is not guaranteed and they can have a hard time with decoding complex error models. BP decodes are also plagued by degeneracy, where distinct physical errors generate the same syndrome and have the exact same logical effect because they differ only by a stabilizer. Standard BP decoding is designed to find the single most likely error string and struggles with this phenomenon which prevents it from detecting distinct physical errors. As a result, BP must be completed with some post-processing techniques like Ordered Statistic Decoding (OSD) [1776, 1777] and other methods [1778].
- **Machine learning syndrome decoding** is another option [1779–1782]. It can use multiple techniques like reinforcement learning [1783], transfer networks [1784], transformer-based neural networks [1785], mixture of experts vision transformers [1786], Hypergraph Neural Networks [1787], and Monte Carlo Tree Search (MCTS) [1788] or Markov Chain Monte Carlo [1789]. One notable work in the domain comes from Google with AlphaQubit 2 whose published benchmarks report decoding faster than 1 μ s per cycle

up to distance 11 for a surface code, on the model version, hardware, batching mode and noise distribution described in that paper, and not on a Google device operating at that distance [1790]. Beyond such distance, training becomes difficult, and [1791] projects that inference would be too slow beyond distance 23 even with low-precision weights such as INT4. These are projections from a stated benchmark setup and not measured limits. It will become a challenge for FTQC systems supporting more than a million gates. An efficient decoder for 12 logical qubits requires 1.90×10^7 parameters and 10 days of training on a Nvidia A100. Scaling could come with the help of attention-based neural decoders although having response time only compatible with slow-gate qubits [1792]. QuEra also tested a machine learning based decoder [1135] which was later improved in 2026 [1793] to support logical error rates of 10^{10} . We can expect progress in the way these systems are trained, and on them running on even more powerful clusters of Nvidia DGX servers or equivalent appliances.

- **Other Decoders.** The list is long, with many decoder proposals that are supposedly more efficient than the above ones. SymBreak is one example, with a better logical error rate at the expense of a moderate time overhead [1794]. SOME (Symmetric One-Hot Matching Elector) is a decoder proposed by researchers in Japan in 2025 that turns a syndrome decoding problem into a QUBO problem that is easy to solve on a classical CPU (so, a different type of QUBO problem than those solving combinatorial optimizations with analog quantum processors) [1795]. The paper claims up to 99.9 $\times$ reduction in variable count (why not 100...?) and ms-level decoding times on a single-threaded CPU, which deserves some scrutiny.
- **Window based decoders** where the decoding problem is split into smaller graphs called windows that are processed sequentially (sliding windows) or in parallel [1796].
- **Special function decoders** like an erasure decoder for qLDPC [1797], a decoder for Tanner codes [1798].
- **Hierarchical decoders** which are stacking several decoders in order to manage the right trade-offs between speed and accuracy [1799].
- A variation of this is proposed with the use of digital annealing, to work with color codes [1800]. Graph neural networks can also decode error syndromes in qLDPC codes [1801, 1802]. A similar approach goes up to using a quantum annealer to run syndrome decoding. This thing becomes unreasonably recursively complicated [1803]! Similar proposals were made in 2025 to operate syndrome decoding with a variational quantum circuits running on the same hardware [1804] and in 2026 using QAOA [1805].

- **Abortive decoding schemes** are envisioned to detect a critical mass of errors and discard an entire circuit shot mid-course to avoid wasting resources [1806].
- **Multistage decoders** combining multiple decoders [1807, 1808].

3D spacetime decoding. In standard QEC with non-constant-time overhead error correction codes, syndrome checks are repeated d time, d being the code distance. It transforms a simple 2D error-tracking task into a massive 3D spacetime computational problem. Traditional decoders like MWPM or BP must navigate this complex 3D graph to distinguish physical errors from measurement false alarms. Knill error correction bypasses this bottleneck entirely by using pre-purified auxiliary states and a single step of logical quantum teleportation, it extracts the syndrome in one high-fidelity shot, flattening the mathematical problem back to 2D so classical decoders can process it in real-time [1809]. But it has a significant space and time cost, requiring additional physical qubits and involving additional routing and connectivity between the qubits. It is mainly relevant for photonic and atomic qubits.

Syndrome decoding classical cost. The classical part of quantum error correction can be significant. It is handled by a classical computer that is as close as possible to the electronics controlling the qubit, whatever the qubit type. A typical configuration would be an FPGA-based system like what IBM is working on, or one or several Nvidia DGX servers using NVQLink, a high-speed interconnect architecture hooking QPUs and DGX servers, with an interface developed jointly with the Department of Energy in the USA [1810]. Depending on the platform, error syndrome decoding can require an HPC [1811] or a simple laptop [1812]. It depends on the syndrome extraction cycle time, error correction code structure and decoder algorithmic complexity.

Google identified a $63\ \mu\text{s}$ latency threshold for syndrome decoding in its Willow 2024 paper [1645]. Fulfilling it will probably need GPUs, particularly for large distance codes [1813]. The turn-around correction cycle must be as fast as possible, under the constraint of physical qubit decoherence time. One key bottleneck here is syndrome decoding. Exact maximum-likelihood decoding is NP-hard in general, but practical decoders solve an easier problem, and the operational constraints are throughput, latency and backlog [1814, 1815]. QEC dynamic code compilation is also a challenge [1816]. Surface code compiling with many logical qubits seems possible in a reasonable time [1530].

Real-time syndrome decoding is a classical bottleneck of quantum computing. It uses the parity errors collected by the stabilizers from an error correction code and solve a complex graph problem to determine where errors must actively be corrected in the logical qubit. The larger the code distance, the more complicated this problem is. With the available current methods, the

decoding speed is usually inversely proportional to its accuracy. The decoding time scales from $O(d^2)$ to $O(d^6)$, d being the code distance, for surface codes decoded using respectively the Union Find heuristic method and the MWPM (minimum weight perfect matching) exact-solution method [1758]. The decoding complexity depends on the specific decoder employed, which in turn depends on the particular QEC code. In the case of color codes, it was demonstrated by Riverlane in March 2026 that exact decoding was an NP-hard problem [1817].

While we can theoretically store data and decode it later (offline decoding), a functional Fault-Tolerant Quantum Computer (FTQC) must decode errors in real-time for multiple reasons. One is the error data backlog problem generated by the staggering volume of data generated by parity error measurements. If the time it takes to process the data and apply the error correction is longer than the whole cycle, errors will accumulate. As a result, this backlog problem is significant for fast qubit types like superconducting and silicon spin qubits. For slow cycle qubit types like cold atoms and trapped ions, error syndrome decoding have more time to process errors. The second reason concerns the deadlines imposed by the circuit. Decoding can be pipelined and its result tracked in a Pauli or Clifford frame instead of being applied immediately, so a delay is acceptable as long as the decoder keeps up with syndrome production and as long as the frame is resolved before the next feed-forward decision, typically a non-Clifford gate or a measurement whose basis depends on it. What must be avoided is an unbounded backlog, not any delay at all. You are also constrained by the coherence time of your qubits and their physical gate speed.

Error syndrome decoding can be very expensive for large distance codes. So far, practical codes implemented in hardware do not exceed distance 7. To execute large quantum chemical simulations, we'll need code distance way above 20 and depending on the platform and its physical qubit operations fidelities and connectivity. Four different things must be separated here: the largest code distance demonstrated in hardware, the largest real-time closed-loop demonstration, the distances reachable in decoder simulation, which go well beyond hardware demonstrations, and the asymptotic scalability of the decoding algorithms themselves. The open problem is real-time decoding at large distance, not decoding at large distance in general. Google developed some machine learning based techniques like with DeepMind AlphaQubit 2 [1790]. It can't exceed distance 11 codes as of the beginning of 2026. IBM prototyped some real-time decoding based on AMD FPGAs in 2025 that was implemented with gross codes of dimension 12 [1818]. Several similar real-time FPGA-based experiments have been implemented [1819], and up to distance 12 codes [1820].

Pre-decoding is a lightweight, preliminary processing step that quickly resolves simple, localized physical errors before passing the remaining sparse error data to

a heavier, full-scale decoder [1821, 1822]. This is the part of decoding that could potentially run at cryogenic temperature for solidstate qubits [1823, 1824].

5.2.9 Logical qubits

With quantum computers available online such as those from IBM having over a hundred qubits, it is the role of software to implement dynamic error correction codes and more precisely, compilers that will transform the developer's code into executable machine code at the physical level of the qubits and integrating QEC code. Given that we have at this point just enough qubits to test small QEC like Steane's 7-qubits codes or small sized surface codes.

Conceptually, a logical qubit sits between a physical qubit (with a small lifetime and prone to significant error rates) and a mathematically perfect qubit (with infinite computing time and zero error). It lasts longer than a physical qubit and should have an error rate that is compatible with the constraints of the quantum circuit size. This acceptable error rate is more or less the inverse of the circuit size, its breadth (number of qubits), multiplied by its depth (number of gate cycles). It should be in the range 10^{-5} to 10^{-15} to enable a wealth of quantum algorithms. The number of physical qubits to be assembled to create a logical qubit first depends on the error rate of the qubits. The higher the qubit error rate, the more qubits must be assembled. This number can reach several thousand qubit physical qubits [1825]. As a result, QEC settings will be algorithm dependent and seem to be bound to be implemented mainly with software and on generic QPU architectures. Beyond the target error rate, the number of logical qubit per physical qubits will also depends on many other factors such as the quantum error correction code used, the underlying physical qubits fidelities and their connectivity as shown in Figure 273.

Old estimates were accounting for a need of around 1,000 physical qubits per logical qubit. This is an outdated view. This number depends on the above factors which are circuit dependent (mainly, its size) and physical qubit dependent (fidelities, connectivity). The larger the circuit, the larger the number will be. The higher the qubit operations fidelities and connectivity, the lower this number will become. It explains why thanks to recent improvements in fidelities and connectivity, the physical overhead of QEC has decreased over time.

IonQ was planning in 2020 to create logical qubits corrected with a **Bacon-Shor** code, a variation of Shor's code with 13 qubits which was the first example of a **subsystem code** that was proposed by David Poulin [1826, 1827]. It was enabled by their good physical qubit fidelities [1828]. For qubits that can be physically well connected with their immediate neighbors, the most often considered error correction is the **surface code**, created between 1998 and 2001. It uses matrices of processing qubits (in white) connected to measuring qubits

(in black) via **Pauli X** (amplitude flip) and **Pauli Z** (phase flip) gates operating on these data qubits as shown in yellow and green. A distance-2 example of this kind should be described from its explicit stabilizer generators, stating separately the number of data qubits, the number of measurement qubits, the boundaries and the code distance. The $[[n, k, d]]$ notation counts data qubits only, so it cannot be used to designate a total hardware count mixing data and measurement qubits. In 2026, they revisited their approach with the walking cats codes [1829].

A surface code with distance d requires d^2 replicated qubits and $d^2 - 1$ measurement qubits, so a total of $2d^2 - 1$ physical qubits to correct a single qubit. d being usually odd, you then have an even number of measurement qubits divided in two equal parts with flip (Z) measurement) and phase (X) measurement qubits as shown below with two example of distance 3 and 5 surface codes.

A commonly used empirical fit for the surface code logical error rate per cycle is $P_L \approx 0.03 (p/p_{th})^{d_e}$. The prefactor and the exponent rule come from a particular code layout, decoder, noise model and definition of a cycle, and the expression should be used only in that regime, with p being the physical qubit error rate, p_{th} being the threshold physical error rate below which logical errors falls with d , and d_e being linked to the surface code distance d with $d_e = d/2$ when d is even, and $= (d + 1)/2$ when d is odd. See Figure 265 and an example in Figure 274 of a surface code with a distance-7 logical qubit created by Google in 2024.




Figure 274: The logical qubit created by Google AI in 2022 assembled a distance-5 surface code with 49 qubits but its fidelity was not as good as the underlying physical qubits. See [1830]. Google planned to create a logical memory qubit with better fidelities than their physical qubits with a distance-7 surface code and 98 qubits. This was achieved in August 2024 (arXiv) and then published in a December 2024 Nature paper as shown in this chart. Note the four qubits at the bottom which are dedicated to leakage error corrections. Source: [1831].

In that 2024 Willow experiment, Google reported below-threshold surface code memories, the distance-7 logical




Figure 273: the various parameters influencing the number of physical qubits to build a logical qubit and to run a whole circuit. The elements in blue are hardware and code inputs and those in green are algorithm inputs. The resources derived from them are co-designed rather than belonging to either group, since code distance, logical error target, factory sizing, routing and decoder requirements all depend jointly on the algorithm and on the architecture. Physical resource estimates depend on the overall circuit size which conditions the overhead of QEC while the number of non-Clifford gates (usually T gates and sometimes Toffoli gates) turns into the cost of magic state factories and gate teleportation. This physical cost can be significant and must be added to the cost of QEC. The total computing time depends on the algorithm size, the QEC overhead, the speed of qubit gates and readout and the compiler/transpiler/optimizer efficiency. The chart also shows the link between QEC resource estimates and power plus energetic costs. Power will mainly depend on the number of physical qubits in the system while energy will depend on this multiplied by the total computing time, including the number of time the circuit must be executed. © Olivier Ezratty, 2023-September 2026.

memory showing a lower logical error rate per cycle than the distance-5 one [1645]. A comparison against a physical qubit is a separate claim that requires naming the physical baseline used. These Willow qubits were in any case much improved compared with Sycamore's original 53-qubit version in 2019. It increased qubit relaxation and dephasing lifetimes through an improved fabrication process and environmental noise reduction near the quantum processor, reduced cross-talk between all physical qubits during parallel operation, reduced drift and improved qubit control fidelity through upgraded custom electronics, implemented faster and higher-fidelity readout and reset operations, reduced calibration errors, improved control electronics calibration, and enhanced dynamical decoupling protocols to protect physical qubits from noise and cross-talk during idling operations.

Constant overhead QEC. A constant rate code is a family of quantum error correction codes where the code rate $R = k/n$ stays above a fixed positive value as the number of physical qubits n increases. The rate is a property of the family and does not refer to any particular algorithm. Whether such a family suffices for a given computation is a separate question, since holding the

total failure probability fixed while the circuit volume grows generally forces the distance, and therefore the overhead, to grow as well. In qLDPC codes, the distance can scale linearly with n , allowing it to keep up with very deep circuits without sacrificing the constant rate.

A constant-space overhead means the ratio of physical qubits to logical qubits stays constant even as you increase the circuit depth T . In standard codes like the surface code, a deeper circuit requires a large distance d to maintain fidelity, which in turn forces n to grow for every logical qubit since n is proportional to the square of d . In contrast, constant-space overhead error correction uses high-rate codes where a single large block can protect many logical qubits with a fixed cost of physical qubits, even for very large circuits. Surface codes and color codes encode one logical qubit per patch with n growing as d^2 , so their space overhead is not constant once the target logical error rate has to be held while the circuit grows. Constant-space overhead is obtained with high-rate families such as qLDPC codes and some concatenated constructions. constant-space overhead is frequently obtained at the expense of computing time cost, and the other way around. Some of these con-

structions come with extremely demanding threshold values in their proofs, far below any experimentally relevant error rate [1832]. Such a number is a probability, and comparing it with Planck’s constant, which is a dimensional physical quantity, has no meaning. Though, constant-space overhead and reasonable time-overhead codes are also proposed [1833–1836].

A constant-time overhead corresponds to a situation where the time needed for one logical error correction cycle, including syndrome extraction and gates, does not grow with the circuit size or depth. In many architectures, increasing the distance d to support deeper circuits slows down individual gates. A true constant time overhead system avoids this slowdown, often through techniques like single-shot error correction or high-degree parallelism in qLDPC decoders.

A constant-spacetime overhead would show-up with systems where the total resource cost (physical qubits and total computing time) is a flat multiplier of the ideal logical space-time volume [1838]. While surface codes exhibit polylogarithmic overhead to handle deeper or wider circuits, constant overhead systems maintain a fixed efficiency. Recent advances in concatenated Hamming codes and quantum LDPC codes have proven that it is possible to achieve constant-space overhead while keeping the time overhead polylogarithmic or even constant, even when accounting for the increased distance required by deep circuits [1833, 1834]. Other results establish impossibility under different assumptions [1839]. These conclusions are not contradictory: they rely on different noise locality and geometry assumptions, gate sets, classical processing models, code families and definitions of overhead, so each theorem has to be read with its own resource metric.

Entanglement purification *aka* entanglement distillation is a technique used in the context of quantum error correction that relies on local operations and classical communication (LOCC). LOCC performs qubit gates and measurement whose classical result is communicated classically to another part of the system that will use it to run some conditional processing. A qubit or gate teleportation are typical LOCC operations that rely on classical communication channels on top of a quantum channel. Entanglement purification is used for example to prepare high-fidelity Bell pairs that are then exploited by various codes like surface codes. These Bell pairs generation can be generated with applying series of CNOT gates on multiple qubits with bad pairs being discarded with post-selection, based on measurement outcomes. It can also rely on stabilizer codes like surface codes or LDPC codes to detect and correct errors in entangled states. It can purify n -qubit GHZ or Bell states. A 2025 study by Microsoft et al. [1840] demonstrated some limits of entanglement purification which may have some implications for distributed quantum computing.

Instruction bandwidth bottleneck is yet another engineering challenge for FTQC and error correction.

Thousands of physical qubits must be driven by software-based quantum error correction. It creates a digital workload from the classical control computer down to the physical qubits and their many auxiliary qubits, in a range exceeding several tens of TB/s just for factoring a 1024 bits integer with Shor’s algorithm! Specific architectures can be designed to handle QEC as close as possible to the physical qubits, ideally in cryo-electronics components and with some microcode sitting at the lowest possible stage in the cryostat (for solid-state qubits), starting at 4K [1841]. Other proposals consist in improving syndrome decoding and corrections parallelism [1842]. The current record for a QEC cycle is about 300 ns to one millisecond which starts to fit the speed of superconducting qubits, which are at this point the most challenging of all [1843].

Hook errors are single-qubit Pauli errors propagating from ancilla qubits to other physical qubits during syndrome extraction. Correcting them usually require additional qubit overhead [1844–1846]. These are therefore not “native qubit errors” but errors related to the way error correction codes are implemented.

Single-Shot Error Correction allow errors to be detected and corrected using only one round of syndrome measurements, even when those measurements are noisy. This is made possible when the code structure ensures that measurement errors do not propagate uncontrollably, a property known as confinement [1706]. Several quantum error-correcting codes support single-shot error correction. These include 4D toric codes, 3D gauge color codes, 3D subsystem toric codes, hypergraph product codes which are qLDPC variants [1847, 1848], quantum Tanner codes, homological product codes, and hyperbolic codes [1612]. These constructions do not share a single mechanism and none of these properties is common to all of them. Single-shot correction requires either redundancy in the syndrome information, through soundness, or confinement under a specified noise model, and these assumptions have to be attributed to each family individually [1706].

How about real implementations of logical qubits? Demonstrations are now numerous, but they cover very different milestones: an encoded state, a repeated memory, break even against a stated physical baseline, below-threshold scaling with distance, logical one- and two-qubit gates, a universal logical gate set, and full algorithmic circuits. Some of these have been reached, including below-threshold surface code memories, and others have not. Table 9 also shows the dominance of post-selection usage, which is not scalable in a real world FTQC scheme.

- **Maryland University** implemented in January 2021 a logical qubit using a Bacon-Shor 13 code with a chain of 15 trapped ytterbium ions that was correcting single qubit errors [1828]. They then used two such logical qubits in a configuration of 32 qubits to implement fault-tolerant 2-qubit gates.

Q Are there perfect logical qubits?

No! All logical qubit platforms exhibit some logical error rate for single and two qubit gates as well as idle states. These errors can't be equal to zero. Quantum error corrections enabling zero errors would have an infinite size in number of physical qubits per logical qubit. It would also be very slow and impractical. Logical qubits and all their error correction machinery are created to be good enough to support a given quantum circuit size for practical purpose.

The closest thing to a noiseless reference is a classical state vector emulation, without tensor network compression. It provides a numerical model of noiseless evolution within finite precision and it is not a logical qubit. Emulators support up to about 20 qubits on a laptop and about 30 to 40 qubits on large server appliances, depending on their available memory which

scales exponentially with the number of supported qubits. They are limited by memory, by runtime and by classical floating point precision, so the number of gates they can apply is not unbounded. Using FP64 (floating-point double precision) provides about 15 to 17 decimal digits of accuracy. Every qubit operation introduces a tiny rounding error, particularly when dealing with small vector state amplitudes. How fast this accumulates depends on the algorithm, its conditioning, the implementation and whether the state norm is renormalized, so no universal gate count can be given. This capability could be extended using FP128 quad-precision, which is supported through compilers on hardware that are usually limited to FP64 precision. At this stage, classical wall clock time already lands in an unacceptable zone with over a year of computing [1837].

- **Quantinuum** created a single Steane color-code also using 10 trapped ion qubits. They created a break-even logical qubit with better fidelities than their underlying physical qubits in August 2022 [1852, 1853]. They extended it in 2023 to three logical qubits used to simulate an H_2 molecule using a Bayesian quantum phase estimation [1855] and an optimized version for a 1-qubit adder [1863]. In June 2024, using a qLDPC code, they created a 4-qubit GHZ state with fidelity between 99.5% and 99.7%, but after postselecting over 98% of outcomes. The uncorrected 4-qubit GHZ state has fidelities between 97.8% and 98.7% [1857]. In September 2024, they went up to create 12-qubit GHZ logical states, again with pre and postselection that they label prerejection and postrejection [1858].
- Another team in **Austria** and **Germany** developed in 2021 a proof of concept of two logical qubits made with trapped ions and using a T gate and magic state distillation [1850]. In 2023, they also implemented a fault-tolerant 7-qubit Steane code with 16 calcium ions [1854].
- A **China** research team implemented a distance 3 surface code using 17 physical qubits ($=3^2+(3^2-1)$) on the 66 qubits Zuchongzhi 2.1 superconducting qubits QPU. It implements repeated error corrections and post-processing error corrections [1851].
- **ETH Zurich** with a team led by Andreas Wallraff did a similar experiment with 17 superconducting qubits in 2021 [1849]. In 2025, it implemented lattice surgery to implement a ZZ logical measurement between two logical qubits of distance 3 with experimental superconducting qubit chips [1864].
- **Google** announced in July 2021 the creation of their first logical qubits with 5 and 21 physical qubits, showing a $\times 100$ improvement in the error rate between the 5 and 21 version [1865]. In 2022, they created a distance 5 surface code logical qubit with 49 qubits ($=5^2+(5^2-1)$) that improved logical qubit error rates as physical qubits per logical qubits grows [1866]. Google researchers indicated then that their logical qubits would be better than their underlying physical qubit starting with a distance 7 surface code, requiring about 100 physical qubits. They reached that point in August 2024 [1831]. We detail this in the part dedicated to Google in the quantum computing hardware section. But this was only a logical memory qubit, not a logical qubit implementing any logical qubit gate correction, let alone a two-qubit logical gate correction.
- A joint **QuTech-Fujitsu-Element Six** team demonstrated in 2022 a fault-tolerant operation of a NV centers based QPU with logical qubits made of 5 physical spin qubits and two additional measurement qubits in a 29-qubit QPU running at 10K [1867].
- A proposal coming from Italy **encodes a logical qubit into a qudit**, a quantum object with many distinct states. On paper, it reduces the number of individual physical objects needed to create logical qubits [1868]. But this theoretical proposal moves the goalpost elsewhere, like making sure that quantum gates have good fidelities and not too much overhead in such settings.

What are the figures of merit of a quantum error correction code and a logical qubit architecture? As shown in Figure 273, a key one is the logical qubit fidelities. You can't claim the creation of a logical qubit without adding its target fidelity. The end goal is to reach between 10^{-9} and 10^{-15} error rates. These rates differ according to the target algorithms. 10^{-9} may be enough

Origin and date	QEC code	# phys. qubits	# log. qubits	Breakeven	Phys. qubit fidelity	Log. qubit fidelity	Gate set
Maryland Univ., 2021 [1828]	Bacon-Shor	15 and 32	2 and 3-GHZ	—	98.5% to 99.3%	99.7%	2-qubit gate
ETH Zurich, 2021 [1849]	surface code	17	1	post-selection	—	—	Pauli
Austria, 2021 [1850]	color code + flags	16	2	yes	—	Bell state @ 75%	CNOT and T
China, 2021 [1851]	surface code	17 of 66	1	no	—	—	CZ
Google, 2022 [1830]	surface code	49 of 72	1	no	—	—	no gates
Quantinuum, 2022 [1852, 1853]	Steane color-code	10	1	yes	2.4×10^{-3} SPAM	1.7×10^{-3} SPAM	preparation
Austria, 2023 [1854]	7-Steane	16	2	—	0.36% (1QB), 2.7% (2QB)	—	CNOT
Quantinuum, 2023 [1855]	error detection	8	3	—	—	—	Clifford
QuEra, Dec 2023 [1856]	color code + qLDPC	256	up to 48	yes, post-selection	99.5% (2QBG)	93% (2QBG)	Clifford and CCZ
Quantinuum, June 2024 [1857]	qLDPC	32	4-qubit GHZ	yes, post-selection	$97.8\% < f < 98.7\%$ GHZ	$99.5\% < f < 99.7\%$ GHZ	Clifford
Quantinuum, Sept 2024 [1858]	tesseract code	56	8 and 12 GHZ	yes, post-selection	—	—	Clifford
Quantinuum & Microsoft, 2024 [1859]	Steane, CSS, and carbon codes	56	1	yes, post-selection	—	“800× improvement on CNOT gates”	CNOT
Google, 2024 [1645]	surface + repeti-code	101/105 and 53/72	1	yes, post-selection	99.67% (2QB gates)	0.14% error (quantum memory)	no gates
Infleqion, 2025 [1860]	many-hypercube code	114	12	yes, post-selection	—	99.48% (2QB gates)	CNOT
Quantinuum, 2026 [1861]	Steane code	98	48	yes, post-selection	—	99.92% (2QB gates)	H, CZ, T
Infleqion, 2026 [1862]	—	80	30	yes, post-selection	—	—	CZ, CCZ

Table 9: Compilation of logical qubit experimented or documented in the few recent years. Table cells are not yet entirely filled. Post-selection results means the error correction is more like an error detection and cannot be used repetitively. The cut-of date for this table is September 2026 with potential missing results. Table 10 contains a more detailed overview of vendors track record per figure of merit in both QEC and FTQC, or in blueprints. © Olivier Ezratty, September 2026, sources in the text.

for condensed matter and Hubbard models simulations, 10^{-15} is required for Shor’s integer factoring for 2048-bit RSA keys while even better error rates are required for complex chemical simulations. The inverse of these error rates corresponds approximatively to the number of T gates to execute in these algorithms, which are the costlier to correct in surface codes. A milestone in this quest is the **TeraQuop** quantum computer, that could execute a trillion quantum operations before a single logical error occurs (10^{12}).

Then comes the overhead with the number of physical qubits per logical qubit but also code time cost,

meaning, how will it slow down quantum computing. A logical qubit must also be able to implement non-Clifford quantum gates in a fault-tolerant manner, an aspect we’ll describe in the next part. There are some documented lower bounds for physical per logical qubit number [1869].

Will logical qubit entirely fit in a monolithic chip? It is far from being obvious when looking at the various vendor roadmaps. IBM plans to scale-out their chips starting with 120 qubits (Nighthawk, introduced in 2025) first with microwave interconnect within very large dilution refrigerators, and then, as late as possible, with

microwave/optical photons transduction and photonic-based interconnect. Trapped ions vendors like IonQ also have/had various scale-out approaches also using photonic interconnect while some are planning to scale their ion traps and assemble them to support enough logical qubits without needing photonic interconnect. Silicon spin qubits could scale to higher levels thanks to being smaller, but their large scale entanglement must be demonstrated like with each and every other platform.

How about applying error correction to qudits and creating logical qudits instead of qubits? Qudits hold more information than qubits and therefore could potentially enable FTQC QPUs with a smaller overhead. This is a research area outside the commercial world at this point [1870].

5.2.10 Magic states

A Shor and Steane code can correct any Pauli error, including Y gate, which is equal to iXZ . It can correct any linear combination of I, X, Y and Z gates with complex numbers. This comes from the fact that any unit operation on a qubit can be expressed as a combination of $IXYZ$ with complex factors: $U = aI + bX + cY + dZ$. This means, indirectly, that these QECs should be able to correct analog and continuous errors such as slight variations in amplitude or phase, e.g., rotations of a few degrees in the Bloch sphere.

Magic states are not there to correct errors. They supply the non-Clifford resource needed for universal fault-tolerant computing, allowing a gate outside the Clifford group such as a T gate (eighth of a phase rotation in the Bloch sphere) to be applied using only fault-tolerant Clifford operations, measurement and a prepared resource state [1320].

Why magic? This is a quantum information science nickname covering the quantum states which bring some exponential quantum advantage, and are outside the Clifford group states. It is also called “nonstabilizerness” because these states cannot be prepared from stabilizer states with Clifford operations alone [1909]. They can perfectly well be encoded in and protected by a stabilizer code, which is precisely what distillation and cultivation do inside a magic state factory. Magic and contextuality are resource-theoretic notions, necessary for universal quantum computing in this framework, and they do not by themselves guarantee an exponential advantage [906, 1910, 1911].

These states are prepared by a process called magic state distillation or cultivation, in a magic state factory, which uses concatenated error correction codes and, usually Steane or Reed-Muller codes [1912–1915]. It has a significant space and time overhead, with a complex distillation pipeline [1916], and with the number of required physical qubits to create a single magic state, of about two orders of magnitude ($\times 100$) compared to

correcting Clifford group quantum gates. It also uses a lot of post-selection [1917].

This overhead tends to decrease thanks to innovations like the “magic state cultivation”, proposed initially by IBM in 2020 [1655], perfected by Craig Gidney in 2024 at Google [1918, 1919], experimented late 2025 also at Google and on Willow in relation to some color code and code switching [1920].

The technique is being improved in multiple domains:

- With using erasure errors and reducing its overhead by researchers in Israel, at AWS and Yale [1921, 1922].
- The MSC-LS protocol proposed by a team in Japan to reduce magic state cultivation space overhead on square-grid connectivity QPUs using lattice surgery in October 2025 [1923].
- A new code family, the RP2 code, which is a variant of the rotated surface code constructed by placing it on a two-dimensional manifold called RP2 [2, 5, 6] proposed in 2025 by Jian-Wei Pan’s team [1924]. It allows the preparation of T states on regular rotated surface codes, which can then be teleported more efficiently via transversal CNOTs, rather than more complex lattice surgery CNOTs. It reduces space-time resources by an order of magnitude with estimates reaching 10^{-9} error rates.
- The generation of magic states cultivation in relation with a surface code logical memory, avoiding the usage of an intermediate color code and code switch, proposed by AWS in 2026 for all-to-all connectivity qubits [1921].

Magic state cost. The magic state factory space-time cost is the product of the number of active qubits for that round, the survival rate at the beginning of that round, and the round length. The overall space-time cost is the sum of the expected space-time cost for each round divided by the total survival rate at the end of the protocol. Until 2024, the resource estimations of FTQC algorithms were often made with counting the number of T gates (“T count”). This overhead could however be significantly reduced using some error detection [1915], with variants of magic state cultivation already mentioned, with some parallelization [1925], with 3D correction codes, which are difficult to implement with actual qubits at this time due to their high long-range connectivity requirements, which can be implemented with some qubit types like with superconducting qubits [1926], with trapped ion qubits, although with limited scaling capabilities, and with cold atoms using atoms displacements in dedicated zones for the execution of two-qubit gates. But, as of 2026, T/Toffoli count remains a key metric for FTQC circuit estimates, particularly for large circuits, since the published cultivation protocols become unfavorable beyond about 10^9 T states under the noise model, target logical

General information			QEC features										FTQC features								
quantum computing hardware vendor	number of logical qubits	supporting physical qubits	main QEC codes	logical memory experiment	break-even logical error rate	leakage error correction	erasure error correction	correlated error correction	< break-even fidelities	single qubit logical gate	single logical error rate	two qubit logical gate	two qubit error rate	lattice surgery	magic state preparation	gate teleportation	code switching	postselection or detection	logical GHZ fidelity	full universal gate set support	
QuEra	96 [1135]	448 [1135]	color codes	[1856]	[1856]	[1135]	[1135]	[1135]	[1856]	[1135, 1856]	[1135]	[1135, 1856]	[1135, 1856]	[1135, 1871]	[1135]	[1135]	[1135]	[1135]	[1135]	[1135]	[1856]
	12 [1860]	114 [1860]	qLDPC variants	[1860]	[1860]	[1860]	[1860]	[1860]		[1860]		[1860]	[1862]		CCZ [1862]	CCZ [1862]		[1862]		[1862]	
Atom Computing	24 [1872]	208 [1872]	Bacon Shor, 4D code, and other codes [1872]	[1872]			[1872, 1873]		[1872]	[1872]		[1872]						[1872]			
Pasqal	2	4																yes			
Quantinuum	48 [1874]	98 [1400]	sc, Tesseract, Iceberg, concat, C ₄ -Helix	[1700, 1858, 1875]	[1875, 1876]		[1858]		[1875, 1876]	[1858, 1859]	[1876]	[1858, 1859, 1875]	[1875, 1876]		[1861]	[1861]		[1858, 1875]	[1875, 1876]	[1861]	
	12 (2026)	256 (2026)	qLDPC bb [1877], walking cats [1829]	[1878, 1879]	[1878, 1879]		[1880]		[1879]		[1829]		[1829]		[1829]	[1829]				[1829]	
AQT	2	30	cc, holographic codes [1881]	[1882, 1883]	[1884]	[1885]				[1883, 1886]	[1883]	[1850, 1886, 1887]	[1888]	[1884, 1886]	[1850]	[1850]	[1850]			[1850]	
Google	1 [1645]	105 [1645]	sc, cc	[1645]	[1645]	[1645]		[1645]	[1645]					[1889]	[1889]		[1889]	no			
IBM	200 (2029)	n×120	gross code (bb)		[1616]	[1616]				[1616, 1890]		[1616, 1890]		[1616]	[1616, 1891–1894]	[1616]	[1895]			[1616, 1894, 1895]	
IQM	4-36 (2026)	54 (2025)	bb, tile codes (qLDPC)	[1618]						[1896]		[1896]									
Alice & Bob	100 (2030))	2000 (2030)	LDPC, elevator codes						[1897]	[1897]		[1897, 1898]		[1899]	[1899]						
PsiQuantum	100 (2030))	1M (2030)	qLDPC, sc, dual rail				[1900, 1901]		[1900, 1901]	[1900, 1901]		[1900, 1901]		[1902]	[1900–1902]						
Xanadu			GKP, qLDPC, bb				[1903–1905]		[1903–1905]	[1903–1905]		[1903–1905]			[1903–1905]						
Quandela	12		Floquet						[1906]	[1907]		[1907]									
ORCA Quantum	12		qLDPC, sc			not applic	[1908]		[1908]	[1908]		[1908]									

Table 10: Inventory of several QEC and FTQC features as experimented or planned by FTQC industry vendors. Cell color codes: green for published experiments, yellow for published experimented work with bad results, blue for blueprints and theoretical papers. Supported physical qubits provides an indication of the physical scale of the experiments. Main QEC codes: main error correcting codes with sc=surface code, cc=color code, bb=bivariate bicycle codes. Breakeven logical error rate: logical qubit demonstrated with a lower error rate than the underlying physical qubits on a single QEC round. Logical gates error rates are irrelevant when the gates are implemented in the compiler. Lattice surgery: for the implementation of CNOT or SWAP gates. Magic state preparation: experimented for the creation of a T state. Gate teleportation: used to teleport a T state on a distant logical qubit. Code switching: technique use to dynamically change a QEC code. Use of post-selection: yes (yellow) and no (green), should be avoided besides magic state distillation since it is not scalable in computing for deep circuits which are the characteristic of FTQC algorithms. Logical GHZ fidelity: implemented with individual logical qubits. This table may be incomplete. (cc) Olivier Ezratty, March-September 2026.

Q Misleading logical qubit labeling

Some vendors label their physical qubits enhanced with embedded error correction or postselection as “logical qubits”. This marketing trend blurs the line between hardware-level noise mitigation and active error correction and FTQC. These “pseudo-logical” or enhanced physical qubits provide measurable fidelity improvements but they are not real logical qubits.

The related cases are autonomous/GKP Qubits which encode discrete logical states into the infinite-dimensional Hilbert space of a continuous-variable oscillator such as a superconducting microwave cavity or an optical mode. This is the case with Nord Quantique (Canada). The qubit natively suppresses certain types of errors like small phase or displacement noise. Some platforms claim autonomous error correction by coupling the oscillator to a dissipative ancilla, stabilizing the GKP state without explicit syndrome measurement and feedback cycles. These

qubits are still physical states living in a hardware component. To achieve true fault tolerance, GKP qubits must themselves be assembled into a higher-level discrete QEC code such as a surface code or color code to handle residual errors, stabilizer measurements, and implement logical gates fault-tolerantly. Without that second layer of active error correction, noise accumulation eventually drives the oscillator out of its code space.

Postselection used by IBM in some papers discards runs where an error is detected, keeping only the successful outcomes. It uses mid-circuit measurements and dynamic circuits, which can check whether a physical or pseudo-logical operation succeeded. If a syndrome indicates an error occurred, that specific shot is thrown away. But postselection is not active error correction, it is error filtering and it doesn’t scale well since its overhead is exponential with the circuit size.

error rate and code distance assumed in those papers, a boundary which is protocol- and version-dependent Figure 275. The reason is that magic state cultivation relies on some post-selection with acceptance rates that decrease quickly as the code distance grows, and annihilates its gains. As a result, magic state cultivation could become a seed for magic state distillation supporting larger circuits, as proposed by Google in 2026 [1927].




Figure 275: Comparison of the cost of magic state generation between cultivation and distillation. While magic state cultivation has a lower space-time cost, it seems limited to circuits under the gigaquop regime. Large circuits requiring more T states will still require the more costly magic state distillation, which may use magic state cultivation for seed magic states to further distill. Source: Figure 3 in [1928].

Other magic states. The most commonplace magic state is $[T|+]_L$, but they are not the only ones. There are proposals to generate other magic states with culti-

vation like $\sqrt{T}|+\rangle_L$, which is the equivalent of half a T gate/state. It could potentially simplify gate synthesis with arbitrary phase angles, which are important in all FTQC circuits, like a QPE or HHL [1929]. qLDPC-based CCZ magic states factories (applying a Toffoli gates) are also proposed by Harvard and tested with cold atoms [1930].

Asynchronous magic state generation. Magic states preparation is designed to be asynchronous with the logical circuit needs. Magic state factories act like asynchronous assembly lines running in parallel to the main data processor. A magic state factory is a specialized area of the quantum computer that spends its time distilling noisy magic states into high-fidelity ones. This process involves many rounds of measurement and “retrying” (post-selection) if an error is detected. Because these retries are random, the factory doesn’t know exactly when a “pure” or “distilled” state will be ready. Once a magic state is distilled, it can be consumed right away or “buffered” in a quiet area of the chip. Then, when the main circuit needs to execute a logical T-gate, it requests a state from the buffer. As long as the factory produces states faster than the circuit consumes them, the computation never has to pause. When the prepared magic state is idle in its logical qubit, it is stabilized with a regular error correction code. The fidelity of a logical T gate will be conditioned by its distillation or cultivation and by the fidelity of its teleportation.

Magic state factory sizing. It is sized to generate multiple T states in parallel inline with the number and frequency of the T gates required in the quantum circuit. It’s linked to the weight of “magic” (also referred to as “contextuality” [906]) in the circuit. All QFT and most modular exponentiations like the one in the Quantum

Q How is a logical T gate practically implemented?

Magic states are actually single qubit quantum states that are outside the space of what you could create with Clifford gate. It corresponds to a purified state that would come from an Hadamard gate followed by a T gate.

This state that can be prepared asynchronously from the rest of the circuit code execution is actually “gate teleported” elsewhere in the quantum register to implement a T gate on an existing qubit and fault-tolerantly. This is called “magic state injection”.

The generated magic state is the equivalent of applying an Hadamard gate on the state $|0\rangle$ which generates a $|+\rangle$ state, and on which a T gate is applied, thus the annotation $T|+\rangle_L$, the L meaning a “logical state” that can be teleported on another logical qubit. The result looks like:

$$T|+\rangle_L = \frac{1}{\sqrt{2}} \left(|0\rangle + e^{i\pi/4} |1\rangle \right)$$

with equal-magnitude amplitudes on both basis states. The state written with unequal trigonometric amplitudes and $\beta = \frac{1}{2} \arccos\left(\frac{1}{\sqrt{3}}\right)$ is a different resource state, of the H type, pointing along another magic direction of the Bloch sphere. Both are used in practice and should not carry the same name.

Magic state injection entangles the magic state with the target qubit using CNOT gates. The magic state qubit is measured in the X basis and then, a corrective Clifford gate (like $S^\dagger$) is applied depending on the measurement outcome. This process teleports the T gate onto the target qubit using only Clifford operations and the prepared magic state. The magic states can be prepared asynchronously vs their needs in computation [1931].

Phase Estimate algorithm used in quantum chemistry and HHL require a lot of T gates which are themselves used to synthesize arbitrary rotation gates along with logical operations. This process can also be handled dynamically according to [1916].

Quantum processor specialization. Different quantum computers could be used for preparing magic states and for consuming them. One proposal from USC and DoE researchers in 2026 uses a fast superconducting based quantum computer for magic state distillation and a cold atom based for computing [1932]. This is a very theoretical proposal that requires optical photons-microwave photons transduction between the two systems. It was also proposed by Q-CTRL in 2026 in its Q-NEXUS architecture based on using a quantum bus and specialized quantum computing units [1933].

5.2.11 Fault-tolerant quantum computing

Fault-tolerance is about extending QEC to support the successful execution of a quantum circuit with logical operations error rates compatible with its size. The first definition of FTQC was crafted by Peter Shor in a seminal 1997 paper [185]. It was the ability to construct a polynomial-sized quantum circuit that can successfully run on a quantum computer despite the presence of decoherence and gate inaccuracies. It comprised multiple elements: the definition of new error tolerance threshold of $O(1/\log^c t)$ inaccuracies and decoherence per gate for a computation with t gates, c being a constant, the use of error correction codes without decoding the data itself, using noisy gates to compute directly on encoded qubits and correcting errors as they occur, correcting decoherence errors and qubit control inaccuracies errors.

Also, QEC codes used repeatedly during long calculations must not introduce more errors than are corrected and should not spread errors in an uncontrollable way to various qubits in the computing register. As Peter Shor recounts: “*To be able to build a quantum computer, it’s not enough to be able to correct errors with noiseless gates; you need to be able to correct errors using noisy gates. This means you have to correct the errors faster than you introduce new ones*” [184]. Formally, what threshold theorems require is that the fault rate per circuit location sit below a threshold, under stated locality and correlation assumptions, and that the decoder keep up with syndrome production and meet its feed-forward deadlines. Speed matters through throughput and deadlines rather than as a race between correction and error creation.

FTQC theoretically allows the execution of algorithms of longer length, whereas without it, we are limited to a few gate cycles. The challenge is to ensure that QEC prevent errors from cascading during the quantum circuit execution.

We must avoid linking one qubit with too many qubits with multi qubit gates in QECs. For this respect, a 7-qubits Steane code is appropriate. And let’s not forget that a CNOT gate propagates flip errors from the control qubit to the target qubit and phase errors from the target to the control. From an operational standpoint, FTQCs creation involves minimizing the number of auxiliary qubits and optimizing the choice of QECs according to the type of errors generated by each type of qubit and quantum gates [1935].

We review here the various components of an FTQC architecture. They are illustrated in Figure 283.

Q What is the difference between QEC and FTQC?

Quantum error correction (QEC) is about stabilizing a logical quantum memory and fault-tolerant quantum computing (FTQC) is enabling logical qubit operations to run a whole circuit correctly.

Quantum error correction protects the encoded information, both while it is idle, which is the “data at rest” case, and during logical operations, since syndrome extraction keeps running while gates are applied. It fights against decoherence and extends the qubit information lifetime. You need as many logical qubits as there are qubits in the circuit you want to run.

Fault-tolerance adds to QEC the following key features:

- **Implementing logical gates** on the logical qubit memory to support a universal gate set, including the support of Clifford gates: X, Y and Z rotations, the Hadamard gate, a CNOT or CZ gate for creating entanglement, and the T gate, which takes the logical gate set outside the Clifford group and enables the synthesis of arbitrary rotation gates. Non-Clifford resources are what

allows a circuit to escape efficient classical stabilizer simulation, which is not the same thing as explaining where an exponential speedup comes from. Some of these logical gates are not directly implemented physically.

- **Avoiding the propagation** of errors within the system due to error correction operations, delays and how the logical gates are implemented.
- **Suppressing the logical error rate** below the level the circuit requires. This calls for a physical fault rate below the threshold of the chosen scheme, a decoder throughput that keeps up with syndrome production, a bounded backlog, and a decoding result available before each feed-forward deadline. Decoding may be pipelined and does not have to be instantaneous.
- **Correct all errors types** that may prevent the successful execution of a quantum circuit like leakage errors, defect errors, loss errors, and correlated errors, which are different across qubit modalities.

See the review in [1934] for more on this.

Transversal gates are implemented with FTQC to avoid propagating errors beyond the corrected qubits. It is an arrangement of links between logical qubits linked together by two-qubit gates. Figure 276 illustrates these links between two logic qubits using a 7 qubit Steane code via CNOT gates. Each of the physical qubits of the logical qubits is connected one by one between the two logical qubits. This is still very theoretical, besides trapped ions, no qubit topology enables this kind of connectivity.

Which gates can be implemented transversally depends on the code. The **Eastin-Knill** no-go theorem states that, under its assumptions, no single finite-dimensional code admits a universal transversal gate set. It does not say that transversal gates are confined to the Clifford group, and some codes do have transversal non-Clifford gates, starting with the $[[15, 1, 3]]$ Reed-Muller code and its transversal T. That’s why we usually need a costly technique named **magic state factory** (based on distillation or cultivation) to generate non-Clifford gates like the T and Toffoli gates. It has a huge cost of two orders of magnitude for physical per logical gates.

The **threshold theorem** demonstrated by Dorit Aharonov and Michael Ben-Or in 1999 stated that it is possible to perform error correction up to an arbitrary desired apparent error rate if the error rate of the single-qubit gates is below a given threshold which is dependent on the error correction code used and the characteristics of the qubits [1733].

The values of 0.1% to 1% quoted for it correspond to specific codes, gadgets and noise models. A threshold constrains the fault rate at every location of the circuit, two-qubit gates, idle steps, preparations and measurements included, and not single-qubit gates alone.




Figure 276: Connecting transversally two logical qubits using CNOT gates. It is easier to implement with qubits supporting many-to-many connectivity. Otherwise these CNOT require a lot of SWAP gates for their implementation.

The consequence of this theorem is to allow the application of error correction codes recursively until reaching the desirable error rate to execute a given algorithm (Figure 271). This is however based on the assumption that qubit fidelities are stable as their number is growing, a feat that is not yet achievable! It also doesn’t correct various sources of errors like isotropic errors affecting simultaneously several qubits [1936]. On the other hand, a variation of the threshold theorem was demonstrated

in 2023 that took account a stable percentage of defects in planar arrays of qubits and includes a QEC protocol for large arrays of defective qubits [1937].

The standard specification in vendor roadmaps for FTQC QPUs is one million physical qubits with 99.9% two-qubit gate fidelity running a surface code QEC. It might enable 100 to 1,000 logical qubits, taking into account the significant overhead of T-gates or Toffoli gates synthesis and correction.

How about the number of logical qubits of a FTQC QPU? There is no universal logical qubit count above which classical simulation becomes impossible, since simulability depends on the circuit structure, the depth, the entanglement generated and the target task rather than on width alone. A meaningful figure comes from a named algorithm with its logical width, its depth, its precision and its classical baseline. The order of magnitude usually quoted combines a few dozen data qubits with a comparable number of ancilla and routing qubits, which is algorithm and implementation dependent, so we end up with about 100 logical qubits and a number of physical qubits that depends on the architecture, ranging from 30 to 10,000 physical qubits per logical qubits. The sheer number of qubits required to build a FTQC awarded it another nickname: Fairy-Tale Quantum Computing!

Most quantum processors won't have more than a couple hundred qubits implemented in a "monolithic" way. Scaling these systems to implement FTQC may require the assembly and interconnection of several QPUs though entangled resources links and the use of the gate teleportation technique. It will drive the creation of specific architectures, particularly for implementing gate teleportation between these qubits using photon entangled resources in the microwave or optical regime [1941].

Logical gates FTQC implementation. Error correction is applied on a quantum memory made of logical qubits. Logical gates are then applied to this logical qubits memory. It must be done "fault-tolerantly" so that it doesn't add more errors on the logical qubits. Errors are corrected at the logical qubit level with logical memory error correction codes. The way logical gates are implemented depends a lot on the QEC-FTQC architecture, on the qubit type and connectivity, and on the type of error correction codes used:

- **CNOTs** can be implemented in multiple ways depending on the hardware platform and error correcting codes. With cold atoms and trapped ions which can support many-to-many connectivity, entangling logical gates like a CNOT can be run transversally between pairs of logical qubits with applying physical CNOT gates to their data qubits. It's sometimes called a TCNOT [1942]. This fault-tolerant operation avoids the propagation of errors to other physical qubits. In qubit systems organized in lattices like with superconducting and silicon spin qubits, CNOT can be implemented using so-called lattice surgery

which merges three separate logical patches into one larger patch for a few cycles and then split them apart again (two data logical patches plus the lattice surgery patch). For longer physical distances between the logical qubits, the intermediate patch is replaced by a "bus" containing series of stabilizer chains for each pairs of boundary data qubits in the code. The logical CNOT is implemented with a series of X and Z stabilizers applied in the patch linking the logical qubit pair. In the Litinski method which is relevant for circuits with a high share of T gates, CNOTs can be absorbed by the final measurements (see Figure 4 in [1943] and [1944]). With qLDPC codes, CNOTs can be implemented transversally, with lattice surgery or with the Litinsky method. There are many other possible methods to implement a logical CNOT like logical gate teleportation, code deformation and braiding, or parity codes as developed by ParityQC in Austria [1945, 1946].

- **Logical X, Y and Z gates** are implemented in two possible ways. The first way is to implement these gates transversally with applying physical X, Y and Z gates to the logical qubit' data qubits. The second and most common way is to implement these gates virtually at the compiler level. It uses the Pauli frame tracking technique which tracks a change of context in the qubit register computational basis within the compiler. This method is implemented by the compiler for blocks of Clifford group gates including the CNOT and SWAP gates. The compiler reorganizes these gates as series of stabilizers measurements. This works until the circuit needs to run a non-Clifford gate like a T gate. At this point the tracked frame has to be resolved, since the accumulated byproduct determines the basis of the next measurement or the classically controlled Clifford correction to apply. This is an adaptive propagation of a classically known byproduct and not a physical flush of accumulated errors. Unknown physical errors are a different object, inferred by the decoder from syndromes, and a Clifford frame is a broader notion than a Pauli frame.
- **Logical Hadamard and S gates** implementation depend on the QEC code. With surface codes, it uses a process involving "code deformation" or "patch deformation" which "rotates" or "folds" the entire code patch by 90° at its edges [1947]. The logical gate is said to be "rotationnally transversal", with a switch between the X and Z basis. With Color codes and concatenated codes, the Hadamard gate can be implemented transversally, meaning with the same physical gate applied simultaneously to all the data qubits in the code. When using surface codes, the S gate is prepared with some state distillation, although it's cheaper than T state distillation.
- **SWAP gate** can be implemented with a series of three consecutive CNOT logical gates but it is very expensive. Alternatives consist in handling these gates in the compiler with changing the labels of the

Q Why QND is key to implement FTQC?

In FTQC, active error correction relies on continuous, mid-circuit measurements of auxiliary qubits to detect and fix errors in real time. What this process requires is that syndrome extraction reveal the stabilizer eigenvalues without measuring the encoded logical information, with controlled back-action and crosstalk on the data qubits. QND readout of the ancillas is convenient but not mandatory: ancillas can be measured destructively and then reinitialized or replaced, provided the schedule tolerates it and the ancillas can be refreshed fast enough. The real difficulty arises when mid-circuit readout disturbs the neighboring data qubits, through crosstalk, scattered light or leakage, or when the measured object cannot be replenished, as with atom loss.

Bypassing this limitation uses deferred measurement, pushing all error syndrome and flag checks to the very end of the experiment [1724]. The error data is

extracted after the operations have concluded. As a result, the system cannot apply dynamic, real-time corrections. Instead, post-selection is used, which involves exploiting the final auxiliary readouts and completely discarding any experimental run where an error flag was triggered. A process that doesn't scale as circuit size get bigger.

Cold atom historically suffered from atom loss during readout, requiring modern workarounds like dual-species setups to avoid post-selection [1938]. Trapped ion face photon crosstalk during ion readout. Ions are shuttled to isolated zones to achieve QND readout [1939].

Even superconducting circuits, which implement QND readout, can rely on post-selection to discard runs where measurement pulses inadvertently leak qubits into non-computational higher energy states [1940].




Figure 277: Example of a surface-code state teleportation using lattice surgery of an arbitrary logical state $|\psi_L^A\rangle$ between two logical qubits using single-qubit and two-qubit measurements, merging and splitting the two logical qubit patches, and measuring the first logical qubit. It is teleported on the logical qubit on the right as $|\psi_L^B\rangle$. The gate teleportation is using a virtual X and Z logical gate. Source: [1884].

logical qubits in the code, or to implement long-range SWAP with a bridge made of stabilizers or logical teleportation. It can also be handled by the compiler when it consolidates a block of Clifford operations.

- **T gates** are usually implemented in two steps. T states, which correspond to running an H and T gate on an initialized qubit, are generated using the process of magic state distillation or cultivation. It's a special zone in the processor that is dedicated to the creation of such a state. It uses noisy T-states and distills them into pure ones. When it's ready, it is teleported using a transversal mechanism to a distant logical qubit, and potentially using some lattice surgery in the case of 2D qubit layouts. Magic state distillation is based on probabilistic measurements

with results that are handled in the Pauli corrections stage along with corrections and Pauli frame change as described in Figure 278. T gates are combined with Clifford group gates like the H gate to synthesize arbitrary rotation gates. Solovay-Kitaev is the historical general theorem to generate arbitrary single-qubit state, while practical fault-tolerant synthesis uses more efficient exact and approximate algorithms such as Ross-Selinger, or protocol-specific resource states, so the synthesis algorithm and the target precision should be quoted with any T-count (see Glossary). These gates are heavily used in circuit components like the Quantum Fourier Transform (QFT).

- **Logical state teleportation** can be implemented in 2D codes with lattice surgery, as in the example in Figure 277.

The circuit compiler plays a key role here. It decides whether Clifford logical gates are implemented physically (transversally or with lattice surgery) or using classical software and the Pauli frame tracking technique with postponed measurement and corrections, based on multiple constraints. The technique is described in Figure 278. It can even move logical qubits around with teleportation instead of relying on lattice surgery patches [1948]. The compiler and optimizer role will be critical with analyzing the user circuit and finding the best QEC and FTQC implementation maximizing performance and minimizing resource consumption.

Logical qubit readout. Measuring a logical qubit at the end of a circuit is performed fault-tolerantly to ensure that the measurement itself does not introduce an uncorrectable error. Logical readout can be implemented in various ways [1549]:

- **Transversal readout:** all physical data qubits within a logical qubit are measured simultaneously in the desired basis, which is usually the Z basis. The logical value is determined by the collective parity of these measurements which is computed classically, not with a quantum parity measurement. This is not simple majority voting. The raw physical outcomes are processed by a classical decoder which uses the history of syndrome measurements to account for potential physical measurement flips.
- **Basis versatility:** fault-tolerant readout is not limited to the Pauli basis (Z , X or Y). By measuring the physical qubits in the X basis or applying transversal gates before measurement, the system can perform logical X or Y readouts. This is essential for implementing non-Clifford gates via injection and for verifying states in teleportation-based architectures.
- **Lattice surgery:** which allows a fault-tolerant projective joint logical measurement. By “stitching” the data qubit patch to an auxiliary qubits patch through a series of multi-qubit parity checks, the logical information is mapped onto the auxiliary qubits. The measured logical observable is projected, so the encoded state is conditioned on the outcome, but the data patch survives as usable encoded hardware instead of being destructively measured qubit by qubit. It also allows joint measurements like $Z_L \otimes Z_L$. This type of logical readout is used mid-circuit, and not at the end of a circuit like the first one [1943].

Lattice surgery is used with surface and color codes to encode and manipulate logical qubits in 2D qubit lattices, and for implementing logical gates between logical qubits, like a CNOT [1950, 1951]. When a logical CNOT is implemented between two logical qubits (*aka* patches) using lattice surgery, a third patch serving as a logical auxiliary is used to mediate this operation. This third patch must have an X-type boundary adjoining

an X-type boundary of one of the two logical qubits, and a Z-type boundary adjoining the Z-type boundary of the second logical qubit. It can have different shapes and sizes: square, rectangle, more or less elongated, L-shaped, U-shaped, and so on. If the two logical qubits are far apart, this patch serves as a quantum bus between the two logical qubits to connect.

The implementation of lattice surgery at a large scale needs to book space between the logical qubit patches in order to be able to create these buses when a CNOT must be performed between two distant logical qubits. Running several CNOTs logical gates in parallel must avoid collisions between the buses, which creates a routing constraint. It could limit the number of simultaneous operations in the quantum circuit. The transpilation layer is in charge of optimizing these routing tasks. Two stages are involved in lattice surgery: merging and splitting. During the merging step, new stabilizers are measured, which cover qubits located along the boundaries of the two patches. This merges the two patches into one, which involves measuring only the new stabilizers. And since the measurement of these new stabilizers is error prone, it must be repeated several times and a correction pass must be made. A similar phenomenon occurs for splitting which also involves measurement operations [1943, 1951]. Due to the significant overhead of lattice surgery, many techniques are proposed to optimize it [1952–1954].

Visualizing lattice surgery in 3D is done with time and space charts showing how X and Z stabilizers are implemented over time and how lattice surgery operates between logical qubits. The Google magic state cultivation paper [1918] shows 3D “pipe diagrams” which present time flows from bottom to top (the vertical axis) and the horizontal planes represent the 2D spatial layout of the physical qubits on the chip. It can also help visualize compiler optimization techniques as shown in Figure 279. Figure 280 shows the evolution over time of T gate synthesis and how magic state cultivation reduces the qubit space and computing time overhead. It compares four different magic-state factory constructions across more than a decade of research. Each panel is a 3D spacetime defect diagram showing how a logical T state is produced inside the surface code. The figure visually demonstrates how the space and time cost of producing a high-quality magic state has shrunk over time. Pipe diagrams can also be used to represent logical operations with color codes [1955]. Figure 281 shows end-to-end cultivation with its components (on the left) and consolidated (on the right).

Code switching is a technique where a logical qubit is “moved” from one type of error-correcting code to another, mid-computation. It dynamically changes how errors are detected by measuring different sets of parity checks named stabilizers. It is used for example to implement a CNOT gate [1959], or a T gate teleportation on a surface code logical qubit [1960, 1961], or for accounting for the variation of fidelities over time




Figure 278: Example of an implementation of a logical circuit with planar error correction codes like the surface code using the Parallel Synthesis Sequential Pauli Computation (PSSPC) method proposed by Microsoft in 2022. The “Pauli frame tracking” method tracks the change of basis resulting from skipped Clifford gate operations in the C_1 , C_2 , and C_3 sub-circuits and the required corrections from non-Clifford rotation outcomes handled in the orange synthesis zone, to maintain the correct phase. The compiler eliminates Clifford gates by commuting them through the circuit and absorbing them into measurements. In this compilation scheme, the remaining explicit spacetime cost comes from non-Clifford operations like T gates, Toffoli gates or arbitrary rotations. Clifford logical operations are real operations too, even when they are absorbed into measurements or into the tracked frame. The “Pauli corrections” blue boxes correspond to the end of a logic block when the Pauli frame buffer is emptied before starting the next block. These blue boxes gather three distinct objects that should not be confused: the classically tracked frame updates coming from commuted Clifford gates, the recovery inferred by the decoder from syndrome data, and the measurement-dependent byproduct of magic state injection. PSSPC delegates non-Clifford gate rotations to auxiliary synthesis qubit units, which can be processed in parallel. The kickback transferring the rotation from the auxiliary qubits to the algorithm qubits remains a sequential process that adds to the total logical circuit computing time. The Azure Quantum Resource Estimator computes the logical depth (total number of logical time steps τ after all Clifford gates have been eliminated and T-states have been synthesized) and the logical width (total number of logical qubits, including algorithm qubits, synthesis qubits, and the massive T-factories). This chart doesn’t show the underlying logical qubit memory continuous error correction. Source: figure 10 from [1949] with comments from Olivier Ezratty, February 2026.

[1962] as well as for implementing trade-offs between speed and accuracy [1963]. Different approaches are adopted with, for example, using a protocol such as logical state teleportation that teleports or transfers a logical state from a code to another code, or by using protocols specific to, or inspired by subsystem codes like by setting gauge operators which converts or distorts a code into another code. Code switching, code deformation, gauge fixing, logical state teleportation and subsystem lattice surgery are related but distinct operations. They differ by their input and output codes, by which checks or gauge operators change, by the logical operation realized, by the measurements involved, and by their fault-tolerance conditions and overheads, and they deserve to be compared on those terms. Code switching using a logical state teleportation protocol is mentioned in QuEra/Harvard’s architecture [1135]. They switch codes between a 3D color code and a 2D color code which allows for the preparation of a magic state in the 3D code which is simpler because it admits

a transverse T gate, and then to teleport it into the 2D code which has the advantage of having a transverse Clifford group, so it only lacks an efficient method of implementing the T gate.

Switching between a surface code and a color code to use the transverse H and S logic gates that are supported by the latter can be achieved by a method named “subsystem lattice surgery” [1964]. The difference between lattice surgery and subsystem lattice surgery is the measurements complexity to connect the logical qubits. Standard lattice surgery performs rigid, multi-qubit parity checks often involving four or more qubits simultaneously along the boundaries of surface code patches to weld them together. Subsystem lattice surgery utilizes subsystem codes, which allow these high-weight parity checks to be broken down into much simpler, lower-weight gauge measurements usually involving only two qubits. This simplifies the hardware implementation as it reduces the number of physical qubits that




Figure 279: various representations of logical operations using the “pipes model” by Google researchers, applicable to surface codes on a 2D qubit layout. It describes the mix of logical qubit error correction in vertical pipes and relations between logical qubits for the implementation of logical gates in horizontal connections: (a) Lattice surgery compilation of a logical circuit containing a projective rotation. Pink boxes correspond to spacetime used for magic state cultivation in the X or Z basis. Grey boxes correspond to waiting time needed for the classical control system to process the results of previous measurements. Green boxes represent a choice of X or Y basis measurements used to apply the phase correction to the injected T state if necessary. The horizontal poles on the bottom and top correspond to CNOT gates. (b) RX gate applied on the system qubit. (c) RZZ gate applied on the two system qubits, represented using the two vertical poles on the right. (d) Standard logical CNOT gate implemented between two logical qubits each adjacent to an auxiliary logical qubit. (e) CNOT logical gate between one stationary logical qubit and another logical qubit moving past it. (f) CZ gate between two logical qubits each adjacent to an auxiliary logical qubit. The yellow corresponds to a transversal logical Hadamard gate. (g) Lattice surgery showing a 2×2 rotation synthesis unit which implements the single qubit rotation V as an alternating sequence of $T \pm 1$, X and $T \pm 1$, Z gates. (h) The rotation synthesis unit consumes a T state every $2d$ surface code cycles. The horizontal poles on the bottom and top correspond to CNOT used to connect the rotation synthesis unit to the system qubit on which the rotation is being applied. (i) Representation of a logical AND gate in a $3 \times 3 \times 5$ region using Z ports. Bottom two ports should be connected to the controls and top port correspond to the target. (j) Toffoli gate implementation which extends the layout of an AND gate from (h) and (i). (k) Another perspective for (j). (l) Hand-optimized pipe diagram showing a compilation of a Mastrovito multiplier. It contains 2 of the 10 layers of the Toffoli gates used in the complete compilation. (m) Hand-optimized compilation of one component of a Karatsuba multiplier designed to occupy a small spatial footprint. (n) Hand-optimized compilation of a Karatsuba multiplier showing 8 of the 51 parity-control Toffoli gates used in the complete compilation. All these operations are aimed at efficiently implementing variants of the DQJ algorithm created by Google researchers in 2024 [1956]. Reading these charts is definitively a new skill. Illustration and (adapted) captions source: [1957].




Figure 280: How the space-time complexity of magic state creation decreased over time, with the work from Google researchers. This partly explains the constant reduction in resource estimates for breaking RSA-2048 keys. Sources: [1918, 1958].




Figure 281: How magic state cultivation is implemented, from bottom to top on the right, with the various components on the left. It consists of “cultivating” the magic state by gradually increasing the code distance ($d = 3$ to $d = 5$, from a triangular color code to a rectangular lattice surface code) while checking for errors. Source: [1918].

must interact at once while still maintaining the same level of fault-tolerant error protection.

Interspersed QEC and logical gates. When implementing a physical logical gate like a transversal CNOT gate, the QEC mechanism is not operating only before and after the logical gate execution, but also, during it [1965, 1966]. The QEC decoder must actively process error syndrome data before, during and after the gate. It generates a spatiotemporal matching graph used by decoders like a MWPM and account for the physical operations implemented during the physical logical gate execution through graph routing techniques. Similar techniques must be implemented during lattice surgery when used to connect multiple logical qubits in a 2D physical layout.

Floor plans. Organizing a quantum processor in logical qubits requires deep engineering. FTQC architectures are usually presented in spare parts, and rarely with the

whole assembled puzzle. The best example was Google Willow with its first logical qubit under the threshold. It was a single memory qubit with error correction, not multiple logical qubits supporting single and two-qubit logical gates. It gets more complicated as multiple logical qubits are spread on a given area in a quantum chip. In the case of Google, these gates were later described in other papers, including for the particular case of efficient preparation of magic states enabling the teleportation of T gates on distant logical qubits. We can hope that at some point, industry vendors will create animated charts describing the dynamics of these real-world systems with qubit initialization, logical operations, magic state preparation and teleportation, and logical read-out. FTQC orchestration, building qubit floor plan, and compiling a logical circuit is a new scientific field in the making [1967–1970]. It involves a lot of complicated space-time trade-offs with regular improvements [1971].

Then, for the algorithms whose input model requires coherent access to classical data, one needs to figure out how to connect these FTQC systems with an error corrected qRAM. This is a conditional requirement, tied to the chosen data access model, and qRAM at computing scale remains speculative. A trade-off will be at stake when designing FTQC systems: how to balance compute vs memory resources. In 2021, Elie Gouzien and Nicolas Sangouard proposed to use only two logical qubits (made with 13,436 physical qubits) and millions of memory qubits to run Shor integer factoring algorithm on an RSA-2048 key, at the expense of a longer computing time of 8 days. This was the extreme case, favoring memory over compute resources [1972]. Another proposal named LSQCA (Load/Store Quantum Computer Architecture) came from RIKEN and NTT in 2025 to organize FTQC systems which improves memory density at a moderate incremental 5% computing cost using small computational space called Computational Registers (CR), space-efficient memory space zones called Scan-Access Memory (SAM) and Magic State Factory (MSF) for generating magic states [1973].

Continuous Variable qubits FTQC is also possible, particularly with CV photonic qubits, mostly using GKP codes [1974, 1975].

Clock synchronization. The orchestration of error correction in FTQC QPUs require some clock synchronization, for example, to make sure that the various parts of a circuit are in the same phase of their syndrome generation cycle [1976]. It also requires very accurate clocks.

FTQC software architecture. Once the overall architecture of QEC and FTQC is understood, it needs to be assembled at the software level. How are all the software components assembled? What is the exact role and structure of the compiler of your logical circuit? Will it compile it entirely at once and generate a physical layout and all controls, or will it be done on the fly, or by chunks of circuits? Where is the optimizing happening? At this point of technology readiness, these are all open questions, with several proposals [1977].

5.2.12 Intermediate FTQC approaches

There are some alternatives to using magic states factories for the support of non-Clifford and arbitrary rotation gates [1978].

One workaround is proposed by **Fujitsu** (Figure 282) with the STAR method, for “space-time efficient analog rotation”. It implements rotation gates *aka* continuous gates natively in analog fashion with precision pulse controls instead of synthesizing them with series of H and T gates to implement variants of the Solvay-Kitaev. Clifford group gates are still implemented with QEC and FTQC. The result is an architecture that scales better at small size regimes and requires much fewer

physical qubits and gates [1979]. Fujitsu estimated in 2024 that they could obtain a quantum advantage with 49,000 physical qubits on a QPE [1980]. This is a paper-specific projection resting on its own assumptions about physical noise, decoder, connectivity and workload, and it has not been reproduced independently. In 2026, they introduced STAR-magic mutation (SMM), a new protocol for generating an arbitrary rotation gate [1981]. They also work on optimizing their compiler [1982].

A variant of the STAR model was implemented by Quantinuum on its Helios system in 2026 [1876], and labeled “pFT” for partial fault-tolerant, and another one was proposed in March 2026 by Qolab [1983]. A derivative was proposed in 2025 by QuEra to operate with a cold atom QPU and support megaqop operations [1871, 1984, 1985]. With that, 10,000 physical qubits could enable 60 logical qubits. The caveat is that it doesn’t scale well beyond that level and that the production of these continuous gates is also probabilistic, requiring some optimizations [1986].

A more radical theoretical pFT proposal came in November 2025 from Liang Jiang’s team in Caltech, with “error-structure-tailored fault tolerance” [1988]. They exploit the actual physical noise structure of the gate, rather than an abstract error model, to build a 1-fault-tolerant continuous rotation gate $R_{Z_L}(\theta)$ that requires any T-gate synthesis. Compared with the Fujitsu STAR and Choi *et al.* partial FT schemes, it could support larger circuits workloads, with about 1.1×10^7 small rotations, with physical errors of 10^{-3} . For running a 50-spin Trotter Heisenberg simulation, the evaluated spacetime cost is $1337.5 \times$ lower than 15-to-1 magic state distillation and $43.6 \times$ lower than magic state cultivation. The rotations are post-selected, with a success probability of about 18.7% at code distance 12. Similarly, a “balanced fusion” technique to generate arbitrary rotation gates was proposed in 2026 by AWS [1989].

Microsoft is also working on early stage FTQC, particularly for the support of Atom Computing cold-atom based systems, one of which being sold to QuNorth in Denmark, and to be ready by mid-2027. It should support at least a hundred logical qubits [1990].

Approximate QEC or Quasi-Exact fault-tolerant Quantum (QEQ) is another partial FTQC approach which sits in-between NISQ and FTQC [1991–1993]. It implements some error correction, but not to the point of creating functional logical qubits [1994]. It still uses some variations of stabilizers and surface codes.

One of these methods named NISQ+ combines aQEC and SFQ driving circuits [1995]. It can help boost the usable circuit volume of NISQ QPUs, which that paper computes as the number of useful qubits multiplied by the number of gates executable under a given error threshold. This product is not Quantum Volume, which is a standardized benchmark built on square random circuits and a heavy output probability test, and the two should not be confused. The related paper raises an




Figure 282: Fujitsu's FTQC proposal reducing the overhead of T gate FTQC correction using arbitrary angle rotation gates. Source: [1987].

interesting point: slow QEC decoders make applications take exponential time to complete, which is kind of problematic! It is explained by the ratio between QEC data generation and QEC data processing is around 2 for syndrome data processing ratio using classical controls. With superconducting SFQ circuits, the ratio is of 0.125 thanks to a very low latency. The proposed SFQ circuit uses a circuit map similar to the qubit circuits topology. It implements an "Approximate SFQ decoder" stabilizer-based algorithm using a union-find algorithm, resets (stopping signal propagation once pairs are found), boundaries (match signals to boundaries) and tie-breaking (chooses single paths among an equal set).

The accuracy threshold of this SFQ circuit is at 5% of physical error rate and significantly interesting at 1%, yielding then a logical qubit error rate of 0.05% with a code distance $d = 9$. It is a much lower required code distance compared to other correction techniques like those using neural networks. The SFQ circuit power consumption is $13 \mu\text{W}$ for a full circuit with a logical depth of 6, has a surface area of 1.3 mm^2 and a latency of 20 ns for QEC. It seems to run at 4K. It could enable the creation of a 78 logical qubits system using 1,000 physical qubits and a computing depth of 4.36×10^6 gates. All these figures come from the assumptions of that single paper and have not been demonstrated in hardware.

5.2.13 Gate-based computing classes

Various gate-based computing classes and denominations are used. They do not all have the same standing: universal quantum computing, NISQ and FTQC are established terms, while several labels below were coined for a single paper, a single vendor roadmap or a single architecture, sometimes overlap, and come with numerical boundaries that are projections rather than accepted definitions. Each is given here with its origin, and the categories are not mutually exclusive:

- **Universal quantum computing** is the quantum computing paradigm in which all quantum algorithms can be implemented from a mathematical standpoint. It must support non-Clifford quantum gates like T
- **NISQ** definition is not really agreed upon. NISQ is best defined by its operating regime, noisy and not fault-tolerant, rather than by a qubit count. John Preskill's 2018 text mentions the 50 to a few hundred qubit range as the scale of interest at the time [1996], which is an indication and not a minimum. A scale was proposed by Dave Bacon at Google in March 2022 [1997], which dealt with a simple scale of number of physical qubits with NISQ being in the thousand, MSQ in the million, GSQ in the billion and TSQ in the tera number of qubits but it didn't hold, since NISQ is powered by quantum error mitigation whose sampling cost grows quickly with circuit volume, which in practice limits useful NISQ execution well before the qubit count becomes the binding constraint. QEC is needed afterwards and we land in the FTQC regime.
- **LSQ** stands for large scale quantum computer and is about having a QPU with a large number of physical qubits. LSQ systems without error correction are useless due to NISQ scaling problems, particularly with variational circuits. Quantum circuit depth is limited by qubit error rates and many quantum algorithms have a breadth (number of qubits) that is smaller than their depth (number of gate cycles). As a result, LSQ will be the physical enabler for FTQC, not large scale NISQ.
- **FTQC** is the main class, which starts with a few logical qubits of average error correction with target error rates of about 10^{-3} to 10^{-4} . We'll maybe have a continuum in the FTQC progress, with logical error rates per operation decreasing progressively, down to about 10^{-15} in the long term, as the number of logical qubits and the circuit volume grow. Targets of that order apply only to very large workloads with their own error budget. These error rates will have to shrink at a faster rate than the increase of logical qubit numbers.
- **eFTQC** stands for early-stage FTQC. It corresponds to the early stages of FTQC with a small number of logical qubits improving qubit fidelities in an intermediate range, in the 50-200 logical qubit range, up to




Figure 283: Consolidation of the error correction and fault-tolerance features needed to build a universal FTQC quantum computer (green for error correction and blue for fault tolerance and logical gate implementation) including the related hardware features (grey) and software features (purple). (*) : without post-selection which doesn't scale for deep circuits needed to run FTQC algorithms. © Olivier Ezratty, September 2025-September 2026.

supporting a MegaQop and down to 10^{-6} logical error rates [1998, 1999]. Researchers are already trying to design FTQC algorithms that would work in such situations, like the randomized Fourier estimation (RFE) which can be used in several quantum machine learning algorithms [2000] and to benchmark it [2001] or variants of a QPE (quantum phase estimate) [1502]. It is also possible that several NISQ variational algorithms that require higher fidelity qubits than are available today could become operational on EFTQC QPUs [2002]. Researchers in Japan are also proposing statistical phase estimation algorithm (SPE) as a replacement for QPE [2003] for the eFTQC phase. Another proposal is about using 3D color codes to solve a sparse IQP (Instantaneous Quantum Polynomial-time) circuit polynomial problem showing a superpolynomial quantum advantage [2004].

- **pFT** or partial fault-tolerance is an architecture where a logical memory is handled by QEC, plus FTQC Clifford logical gates, but non-Clifford states/gates are implemented in a NISQy fashion. This is part of the eFTQC trend. The technique proposed by Fujitsu and named STAR, is also suggested by Quantinuum and Qolab.
- **Practical FTQC** is about FTQC providing a generic quantum advantage for industry grade applications. It would require at least a couple hundred logical qubits and the support for a gigaqop.
- **FASQ** for fault-tolerant application-scale is yet another acronym invented by John Preskill along with Jens Eisert in October 2025, which means “useful FTQC” [2005].
- **PISQ** for Perfect Intermediate Scale Quantum is a proposal from QuTech scientists that corresponds to the arrival of 50+ FTQC “perfect” qubits QPUs. They advocate to get ready for it in parallel with all the efforts related to NISQ systems [2006]. Basically, it is about creating quantum algorithms by using classical emulation solutions in a higher regime than mostly used nowadays, above 40 qubits.
- **VLSQ**: this is large scale FTQC with several orders of magnitude larger number of logical qubits used to run chemical simulations, Shor integer factoring and large scale optimization and industry scale algorithms.
- **FTDQC** is a new term, meaning “fault tolerant distributed quantum computation”, which could potentially be implemented with long distance quantum communication, even with satellites [2007].

5.2.14 FTQC resource estimations

FTQC and QEC generates both space and time overheads compared to circuit execution on NISQ system.

This is the price to pay to run larger quantum circuits. As a result FTQC resource estimations play a key role at multiple levels: to know the system size required to run specific quantum algorithms, to gauge their viability over time in comparison with physical qubit number scaling and industry vendor roadmaps, and then, to understand and adjust any space and time trade-offs.

We deal with resource estimates tools in general in the software tools section of this book. Most of these tools have been designed for FTQC architecture rather than NISQ ones. Their inputs are the circuit size, sometimes weight of non-Clifford gates, and then physical qubits fidelities, gate times and error decoding cycle times. See a 2026 paper on this topic in [2008].

We will look here are how to compute and evaluate FTQC space and time overheads, which are their key drivers, and their potential trade-offs. Then, we will finish with looking at FTQC application needs.

FTQC space overhead.

There are at least three physical overheads linked to the implementation of a full FTQC architecture [1943, 1950, 2009, 2010]:

- **QEC overhead.** This is linked to the encoding rate of the various QEC codes used. The number of physical qubits to consider should always include all required ancilla qubits. This concerns all qubit types. The overhead is larger for low-fidelity physical qubits and low connectivity qubits. With surface codes, one logical qubit occupies about d^2 data qubits, so the overhead per logical qubit grows as $O(d^2)$. The qLDPC gain is better expressed through block parameters: a family with $n/k = O(1)$ keeps a constant number of physical qubits per logical qubit as the block grows, provided the distance grows fast enough with n to hold the target logical failure probability. Constant rate does not mean constant total resources as reliability requirements increase, and check weight and connectivity must be quoted alongside.
- **Lattice surgery overhead.** This overhead is quite variable. When connected logical qubits are just nearby, the overhead is minimal. When they are distant and need patches and buses, the space overhead scales with the routing distance L and code distance d as $O(L \times d)$, frequently representing an additional $O(d^2)$ footprint per long-range connection. Since resolving the boundaries fault-tolerantly requires $O(d)$ measurement cycles, the dynamic space-time volume of these operations scales as $O(d^3)$ [2011]. As illustrated in Figure 268, lattice surgery adds a significant physical overhead on the FTQC system. In the layouts studied in these papers, the routing space reserved for it can approach the footprint of the logical qubits themselves. That ratio belongs to a given layout, routing policy and schedule and is not a general property. This overhead is to consider mainly for architectures restricted to 2D nearest-neighbor connectivity, whereas modalities enabling many-to-many

connectivity and transversal gates, like cold atoms, can bypass much of it at the cost of a significant time overhead due to atom shuttling and slower physical gates.




Figure 284: An example of low-overhead lattice surgery between two neighbor distance-3 surface code logical qubits. The greyed zone corresponds to the surgery connecting the two logical qubits with $O(d)$ additional physical qubits (7 physical qubits here). Source: [2012].

- **Magic state factory overhead.** This depends on the share of non-Clifford gates (e.g., T or CCZ/Toffoli gates, and in the first place, the number of arbitrary rotation gates $R_z(\theta)$ to support in the logical circuit. This deals with all qubit types. The spatial footprint of a single distillation block such as a typical 15-to-1 factory scales as $O(d^2)$, similarly to a surface code. To suppress the logical error rate ϵ to algorithmic requirements, multiple levels of distillation are often required, leading the total space overhead to scale roughly as $O(d^2 \log(1/\epsilon))$. In several published large-scale resource estimates, parallelization demands result in these factories consuming a large share of the physical qubit budget, in the 50% to 90% range. That share depends on the algorithm's non-Clifford count, the factory protocol, the target logical error rate and the scheduling, and is not a constant. Their sizing must ensure that the logical circuit always has sufficient magic states prepared for its needs [2013]. Compilers may also optimize magic state generation to avoid peak demand [2014].

FTQC time overhead

FTQC time overhead depends on the used error correction codes and on their overhead but also on how magic states are generated and teleported on logical qubits. For the surface code and magic state distillation architectures found in the literature, the computing time overhead usually sits between one and two orders of magnitude, a range that depends on the code distance, the schedule, the routing and the factory throughput of each specific architecture. To give an example with superconducting qubits, an average physical gate can last about 50 ns while a full error code cycle can be 1 μ s.

A full appreciation of computing times across multiple qubit modalities requires adopting a holistic view.

Some qubits operations may be quite slow, like with cold atoms and trapped ions, but this can be compensated by many-to-many qubit connectivity and higher physical gate fidelities for trapped ions which enable lower encoding rate for error correction, and faster error syndrome decoding than gate times. Some papers show that trapped-ions based FTQC could be as fast as superconducting qubit FTQC thanks to using constant-time overhead for both error correction and transversal gates implementation, benefiting from a better qubit connectivity, starting with a code distance of 30 [2015].

Objective comparisons between various future FTQC platforms will have to be done using actual standard algorithms like a QPE (quantum phase estimate) for some chemical simulation or Shor integer factoring.

Quantum software developers can use the various resource estimators developed by academics or industry vendors. They are designed to estimate the quantum computing hardware resources and computing times using as inputs either the algorithm itself or an estimate of its number of gates, the error correction code used, the physical qubit quality and qubit gate and readout times. Among these are the Azure Quantum Resource Estimator that Microsoft launched in November 2022 [1949] and which is open-source since January 2024 [2016]. Google launched QUALTRAN in 2023. It also simulates errors propagation [2017]. PsiQuantum launched Bartiq in May 2024, in connection with their QREF (quantum resource estimator format) [2018]. Other simulators are designed to really account for all QEC/FTQC detailed processes [2019].

The base unit of time in a surface code is the logical code cycle (t_{cycle}), the time needed to perform the syndrome extraction rounds associated with one logical step. As an order of magnitude:

$$t_{cycle} \approx r \times (t_{gates} + t_{readout} + t_{reset})$$

with r the number of syndrome extraction rounds, often taken equal to the code distance d for a surface code memory and equal to 1 for single-shot codes, t_{gates} being the time to execute the 4 to 6 CNOT gates of a parity check, which are applied in parallel layers across the patch, $t_{readout}$ being the time to read out the auxiliary qubits and t_{reset} their reset time, these last two often overlapping in practice. A logical operation does not necessarily take exactly d rounds. We usually don't account for the time to implement the single qubit gates to correct errors on data qubits since it is usually tracked by the compiler.

Decoding does not enter this expression as an additive term, since it can be pipelined and run concurrently with the following rounds. What it must satisfy is a throughput condition, so that the decoder consumes syndromes at least as fast as they are produced and the backlog stays bounded, and a deadline condition, so that its result is available before the next feed-forward decision, typically a magic state injection correction or

Q Has a real full-stack FTQC architecture been demonstrated?

As of September 2026, not yet. It would correspond to the last column in Table 10 which is currently in yellow. The current logical circuits experiments generate logical gates and execute a logical circuits but each time, with some limitation.

Rather than ranking platforms, it is more informative to compare them milestone by milestone: repeated QEC cycles, below-threshold scaling, logical one- and two-qubit gates, a universal logical resource including non-Clifford states, real-time decoding, the post-selection policy used, and full algorithm demonstrations. One of the most complete demonstrations as of September 2026 came from Quantinuum [1861], but without generating a logical T state using distillation. They teleported a regularly prepared T state on

a logical qubit. They used post-selection but only upstream using a Repeat Until Success (RUS) method when generating this T state and a $|0\rangle$ state. They executed small QAOA and HHL logical circuits.

Then, in their 2025 experiment, QuEra [1135] used only 200 atoms from their 448 atoms setup. They tested logical CZ gates (using both lattice surgery and gate teleportation, for the sake of comparing them), arbitrary rotation states teleportation, and CCZ (Toffoli) gates. But they use atom-loss related post-selection which doesn't scale at all.

In 2022, AQT [1850] tested all logical qubit gates but used post-selection through the use of flag qubits, and with 17 ions.

a measurement whose basis depends on it. A realistic timing model is therefore a critical path through a named architecture, separating the physical round, the decoding pipeline, the logical cycle, the feed-forward deadline and the magic state factory bottleneck.

The total computing time will look like:

$$T_{total} \approx D \times t_{cycle} \times \Phi$$

with D being the circuit depth and Φ a multiplier representing the bottleneck of T-state production and implementation of logical gates. The multiplier coming from the T-factory can be only one if the distillation factories can cope with the circuit needs, avoiding the wait for magic states.

The green part in Figure 273 shows that the total computing time and the physical cost is driven by your circuit size. FTQC systems will need to dynamically adjust their logical qubits code distance and physical layout based on the circuit size, which will influence the number of physical qubits per logical qubit. It may lead to specific optimization techniques to avoid wasting resources on the FTQC system.

Let's consider two opposite cases with one circuit containing a million operations and another one with a trillion operations. If you run the first circuit with a FTQC configuration designed for the second circuit, your code will be slower, and the physical resources and energy will probably be bigger. We may argue that the energy consumed may be the same whether the FTQC system physical qubits are used or not. But at least, computing time won't be the same since larger distance error correction code may be slower to run, unless they have a so-called constant-time overhead.

We can also consider how end-users and cloud vendors will run FTQC operations. For an on-premise FTQC system dedicated to a narrow range of use cases, the

resource allocation may be static since it may be predetermined. In the case of cloud vendors running multiple FTQC systems and allocating their resources dynamically based on multiple user requirements, resource allocations will have to be dynamic regarding QEC code distance, number of quantum processors used, how they are interconnected, with side effects on code compilation tasks due to the need to implement circuit partitioning.

FTQC space-time trade-offs

In 2026, many papers we published with updated resource estimates to break either RSA-2048 or ECC-256 keys. These papers were based on FTQC architecture and various assumptions. All were using simple noise models and 99.9% qubit fidelities at scale, and used monolithic processor architectures (at the exception of Q-CTRL in [1933] and U. Chicago 2026 in [2020]).

Figure 285 and Figure 286 exhibit the various space-time trade-offs in the resource estimates.

The colored zones in Figure 285 flag regions that look unrealistic from an engineering standpoint, and the estimates they contain come from heterogeneous architectures which differ in logical circuit, arithmetic, code, noise model, connectivity, factory design, parallelism and classical control. Computing times exceeding several months are too long when considering the MTBF (mean time between failure) or some critical components in cold-atom quantum computers, like high-power lasers given each processor need over half a dozen lasers. The too-many qubit zone starts with about 2 million physical qubits which is beyond the reasonable zone for monolithic qubit processors. In that regime, for both superconducting and cold atom qubits, some quantum processor interconnect technology will be required. They are accounted for in the U. Chicago estimate [2020] but with an unreasonable computing time reaching about a year.

Q Can NISQ applications run on FTQC systems?

One expectation seen in many papers is that some limited NISQ applications could deliver their promise with FTQC systems thanks to their support of larger circuits.

This may be true in some situations but not in general, and particularly for variational circuits (VQE, QAOA). These circuits are plagued by various shortcomings like the scaling of the number of circuit shots (VQE) and the barren plateau effect (both). The scaling of the classical optimizer that fine tunes the weights of the arbitrary rotation gates in ansatzes is not obvious as well.

FTQC applications get rid of these constraints by adding their own. Costly classical optimizing cycles

are replaced by costly data encoding techniques. The larger the data set, the costlier is this process, which is why large scale quantum machine learning on classical training data faces a data loading cost that often dominates, particularly to run inferences. Ground state problems then tend to move from VQE toward QPE and its variants, and optimization problems toward other formulations. These are not one-to-one replacements: whether the fault-tolerant version outperforms the variational one, or a classical method, depends on the problem formulation, the input model, the required precision, the state preparation cost and the classical baseline.

The green zone seems the most reasonable one, although still challenging, given the green Pinnacle favorable regression is based on some unrealistic qubit connectivity of 1-to-10, down to 1-to-8 in September 2026 [1638].

The two regressions in the ECC chart in Figure 286 are curiously aligned, which deserve some scrutiny although a new point coming from IonQ in September 2026 shows more directly the effect of slow gate and readout operations.

FTQC application needs

FTQC application needs are described in “use cases” like the ones we describe later in the book in the quantum computing applications section, and flag with an **FTQC** badge.

A bit like for breaking RSA-2048 or ECC-256 keys like mentioned before using Shor’s integer factoring or dlog algorithms, good estimates indicate the circuit size in number of logical qubits and circuit depth which helps compute the logical memory fidelity requirement and QEC overhead. Then, it adds a number of non-Clifford gates (T or Toffoli) which will drive the cost of the magic state factory and magic states teleportation. Based on hardware specifications, one can then compute the physical resources in physical qubits and computing time, to run the application. A good resource estimate also provides indications on the classical computing costs of the solution, which is rarely provided.

In the “The Grand Challenge of Quantum Applications” November 2025 paper, Google researchers laid out an clear FTQC path covering five stages, with the (needed) discovery of new algorithms in abstract settings (Stage I), the identification of problem instances bringing some computational advantage (Stage II), establishing such an advantage for real world applications (Stage III), full-stack resource estimations (Stage IV) and application deployment (Stage V) [2026]. On stages II to IV, they provided resource estimates for known algorithms. In Stage III for example, they compiled the needs in non-Clifford gates for solving various chemical simulation problems (Figure 287). The numbers show that most application require way beyond the meqaquop planned by some hardware vendors by 2030. Resource estimates beyond 10^{12} Toffoli gates correspond to FTQC systems that, extrapolating from current vendor roadmaps, would not be available before 2035. That date is an extrapolation from projections and not an observation. You can also notice the relatively moderate number of simulated orbitals in the estimates. We explain in the algorithm section of this book why these numbers are limited.

See also Figure 3 in [2027] which makes an inventory of key algorithms and their number of Toffoli gates and logical qubits needs.




Figure 285: Space-time resource estimates to break RSA-2048 including the latest 2026 proposed FTQC architectures. The chart is using a double logarithmic scale with the number of physical qubits in X and Shor factoring computing time in Y. These resource estimates all rely on several assumptions, such as simplistic qubit noise models with 99.9% fidelities for all operations (single and two-qubit gates, readout and idle state). The chart exhibits several zones: in orange, where computing time is prohibitive in months and years, sometimes even beyond the MTBF (mean-time between failure) or key components like laser, in pink with too many qubits (above 2 million) which can be positioned far in the future at best, and finally in green, a relatively reasonable space with computing times below two months and a number of physical qubits between 50K and 2M. This zone is seemingly accessible only to solid-state qubits with fast gate operations. You notice the linear regression for specific qubit types (atoms, superconducting) and error correction types. The Pinnacle steep regression for superconducting qubits may be due to the qubit connectivity assumption. Data sources: Gidney 2019 [2021], Gidney 2025 [1958], Q-CTRL 2026 [1933], Iceberg Quantum Pinnacle [1634, 1638], U. Chicago 2026 [2020], Oratomic 2026 [2022], QuEra 2025 [2023]. © Olivier Ezratty, September 2026.





Figure 286: Space-time resource estimates to break ECC-based 256 bit keys with the latest 2026 proposed FTQC architectures. Like for breaking RSA-2048 keys, they all rely on several assumptions, such as simplistic qubit noise models with 99.9% fidelities (for cold atoms and superconducting qubits) and 99.99% for two-qubit gate fidelities in the IonQ paper, and qubit connectivity depending on the cases. There is also a difference in the appreciation of the wall clock in the papers. The IonQ paper computes 25.7 days with a 63% chance of success. The chance of success is boiled in the Alice & Bob estimate but not in the other papers. Shor discrete log algorithm physical resources and computing times are lower than for breaking RSA for regular key sizes. Data sources: Oratomic 2026 [2022], Alice & Bob repetition cat-codes 2023 [2024], Google Babbush 2026 [1635], IonQ 2026 [2025]. © Olivier Ezratty, September 2026.

Horizon	Model problem	Stage III applications	Problem size	Quantum advantage	Stage IV(b)		Stage IV(c)		
					Logical q.	Toffolis	Phys. q.	Time	
Nearest	Condensed matter physics models	2D Ising quench†	121 sites	Evaluated vs. MPS [96]	140 [162]	$4 \cdot 10^4$	$6 \cdot 10^4$	1 second	
		2D Hubbard quench†	128 spin orbitals		160 [163]	$1 \cdot 10^6$	–	–	
		Heisenberg quench†	70 sites	Evaluated vs. exact sim.	131 [164]	$6 \cdot 10^6$	–	–	
		3D uniform electron gas ground state	64 electrons 64 orbitals	Evaluated vs. QMC [165]	132 [166]	$4 \cdot 10^7$	$3 \cdot 10^5$	4 hours	
		2D uniform electron gas ground state	49 electrons 64 orbitals		140 [167]	$7 \cdot 10^8$	–	–	
Medium		NMR spectroscopy†	32 spins	Plausible vs. TEBD [168]	100 [168]	$3 \cdot 10^7$	–	–	
		Sachdev–Kitaev–Ye†	200 sites	Plausible but not assessed	300 [169]	$3 \cdot 10^7$	–	–	
		Chemical dynamics (2nd quantization + vibrations)	Singlet-fission organic solar cells† [170]	4 states 11 modes	Evaluated vs. MCTDH [171]	1100 [170]	$3 \cdot 10^7$	–	–
		Nuclear dynamics (1st quantization)	Pionless effective field theory†	40 nucleons	Plausible but not assessed	500 [172]	$6 \cdot 10^8$	–	–
		Dynamic structure factor evaluation (2nd quantization)	Battery redox reaction by Li_2MnO_3 (10^4 shots) [173]	18 orbitals	Plausible beyond RAS methods [174]	100 [173]	$9 \cdot 10^{11}$	–	–
Furthest	Chemical ground states (2nd quantization)	Carbon fixation by Ru-complex [41]	56 orbitals	Not assessed	900 [43]	$2 \cdot 10^8$	–	–	
		Oxidation by p450 enzyme [127]	58 orbitals	Evaluated vs. DMRG [127]	1200 [43]	$5 \cdot 10^8$	–	–	
		Nitrogen fixation by FeMo-cofactor [137]	76 orbitals	Evaluated vs. DMRG [130]	1500 [43]	$1 \cdot 10^9$	$5 \cdot 10^6$	9 hours	
		Structure of LiPF_6 battery electrolyte [175]	116 orbitals	Plausible but not assessed	18000 [175]	$9 \cdot 10^{11}$	–	–	
		Bridged Mo_2 [176]	70 orbitals	Evaluated vs. DMRG [176]	4100 [176]	$3 \cdot 10^{12}$	$6 \cdot 10^7$	2 years	
	Chemical dynamics (1st quantization)	Reaction of NH_3 and BF_3 [177]	40 electrons	Plausible but not assessed	800 [177]	$2 \cdot 10^{15}$	–	–	
		Stopping power for inertial fusion [37]	218 electrons	Plausible beyond TDDFT	6200 [37]	$1 \cdot 10^{15}$	–	–	
	Material ground states (1st quantization)	Structure of LiNiO_2 battery cathode [38]	92 electrons 58 orbitals	Plausible beyond CCSD	1400 [40]	$1 \cdot 10^{14}$	–	–	
	Material ground states (2nd quantization)	Heterogeneous NiO / PdO catalysis [178]	[2, 2, 1] supercell		80000 [38]	$1 \cdot 10^{12}$	$9 \cdot 10^7$	9 months	
			64–72 atom supercell	Plausible but not assessed	100000 [178]	$3 \cdot 10^{11}$	$2 \cdot 10^8$	3 months	

Figure 287: Compilation of Stage III FTQC applications in quantum chemistry with their basic algorithm resource estimates corresponding to Stage IV, with a few numbers on physical qubits and computing time. Source: Google in [2026].



5.2.15 Bibliography and notes

These are the bibliographical references and notes for the **Error handling** subsection.

- [184] [The Early Days of Quantum Computation](#), by Peter W. Shor, arXiv, October 2022 (10 pages).
- [185] [Fault-tolerant quantum computation](#), by Peter W. Shor, arXiv, March 1997 (11 pages).
- [906] [Contextuality of Quantum Error-Correcting Codes](#), by Derek Khu, Kishor Bharti, et al., PRX Quantum, January 2026 (29 pages).
- [1135] [A fault-tolerant neutral-atom architecture for universal quantum computation](#), by Dolev Bluvstein, Mikhail D. Lukin, et al., Nature, November 2025 (30 pages).
- [1181] [Quantum computations: algorithms and error correction](#), by Alexei Kitaev, 1997 (60 pages).
- [1320] [Universal quantum computation with ideal Clifford gates and noisy ancillas](#), by Sergey Bravyi and Alexei Kitaev, Physical Review A, February 2005 (14 pages)  (preprint on [arXiv](#)).
- [1325] [Fast, accurate, high-resolution simulation of large-scale Fermi-Hubbard models on a digital quantum processor](#), by Gavin S. Hartnett, Khadijeh Sona Najafi, Aleksei Khindanov, Haoran Liao, Michael Schutzman, Michael R. Hush, Michael J. Biercuk, and Yuval Baum, arXiv, May 2026 (35 pages).
- [1329] [Optimal Qubit Reuse for Near-Term Quantum Computers](#), by Sebastian Brandhofer, Ilia Polian, and Kevin Krsulich, arXiv, July 2023 (11 pages).
- [1330] [99.9904% SPAM Fidelity with barium-137 sets the standard and creates a further step towards solving some of the world's most intractable problems](#), by Kortny Rolston-Duce, March 2022.
- [1331] [Dancing with qubits](#), by Robert Sutor, 2019 (516 pages).
- [1332] [Practical Introduction to Benchmarking and Characterization of Quantum Computers](#), by Akel Hashim, Long B. Nguyen, Noah Goss, Brian Marinelli, Ravi K. Naik, Trevor Chistolini, Jordan Hines, J. P. Marceaux, Yosep Kim, Pranav Gokhale, Teague Tomesh, Senrui Chen, Liang Jiang, Samuele Ferracin, Kenneth Rudinger, Timothy Proctor, Kevin C. Young, Irfan Siddiqi, and Robin Blume-Kohout, PRX Quantum 6, 030202, August 2025.
- [1333] [Fault tolerance against amplitude-damping noise using Bacon-Shor codes](#), by Long D. H. My, Akshaya Jayashankar, Prabha Mandayam, and Hui Khoon Ng, arXiv, March 2025 (24 pages).
- [1334] [Leakage Benchmarking for Universal Gate Sets](#), by Bujiao Wu, Jianxin Chen, et al., arXiv, April 2024 (27 pages).
- [1335] [Randomized Benchmarking with Leakage Errors](#), by Yi-Hsiang Chen and Charles H. Baldwin, arXiv, July 2025 (20 pages).
- [1336] [How does the DRAG \(Derivative Removal by Adiabatic Gate\) technique help mitigate the Stark shift and avoid unwanted transitions in transmon qubits?](#), by EITCA Academy, June 2024.
- [1337] [Counteracting systems of diabaticities using DRAG controls: The status after 10 years](#), by L.S. Theis, F. Motzoi, S. Machnes, and F.K. Wilhelm, arXiv, September 2018 (7 pages).
- [1338] [Ultrafast Single Qubit Gates through Multi-Photon Transition Removal](#), by Y. Gao, A. Galicia, J. D. Da Costa Jesus, Y. Liu, Y. Haddad, D. A. Volkov, J. R. Guimarães, H. Bhardwaj, M. Jerger, M. Neis, B. Li, F. A. Cárdenas-López, F. Motzoi, P. A. Bushev, and R. Barends, arXiv, November 2025 (14 pages).
- [1339] [Removing leakage-induced correlated errors in superconducting quantum error correction](#), by Matt McEwen, Rami Barends, et al., Nature Communications, 2021.
- [1340] [Overcoming leakage in scalable quantum error correction](#), by Kevin C. Miao, Matt McEwen, et al., arXiv, November 2022 (7 pages).
- [1341] [Leakage Mobility in Superconducting Qubits as a Leakage Reduction Unit](#), by Joan Camps, Mark Turner, et al., arXiv, June 2024 (4 pages).
- [1342] [Improved error correction with leakage reduction units built into qubit measurement in a superconducting quantum processor](#), by Yuejie Xin, Sean L. M. van der Meer, Marc Serra-Peralta, Tim H. F. Vroomans, Matvey Finkel, Hendrik M. Veen, Mark W. Beekman, and Leonardo DiCarlo, arXiv, November 2025 (18 pages).
- [1343] [Efficient and Scalable Architectures for Multi-Level Superconducting Qubit Readout](#), by Chaithanya Naik Mude, Satvik Maurya, Benjamin Lienhard, and Swamit Tannu, arXiv, March 2025 (7 pages).
- [1344] [Coupler-Assisted Leakage Reduction for Scalable Quantum Error Correction with Superconducting Qubits](#), by Xiaohan Yang, Dapeng Yu, et al., arXiv, October 2024 (25 pages).
- [1345] [Fast Unconditional Reset and Leakage Reduction of a Tunable Superconducting Qubit via an Engineered Dissipative Bath](#), by Gihwan Kim, Oskar Painter, et al., arXiv, November 2024 (17 pages).
- [1346] [Can Quantum Computers Do Nothing?](#), by Alexander Nico-Katz, Nathan Keenan, and John Goold, arXiv, June 2024 (11 pages).
- [1347] [Linear optical quantum computing with photonic qubits](#), by Pieter Kok, William J. Munro, Kae Nemoto, Timothy C. Ralph, Jonathan P. Dowling, and Gerard J. Milburn, Reviews of Modern Physics, January 2007 (40 pages)  (preprint on [arXiv](#)).
- [1348] [Error models for mode mismatch in linear optics quantum computing](#), by Peter P. Rohde and Timothy C. Ralph, Physical Review A, June 2006 (11 pages)  (preprint on [arXiv](#)).
- [1349] [Loss tolerance in one-way quantum computation via counterfactual error correction](#), by Michael Varnava, Daniel E. Browne, and Terry Rudolph, Physical Review Letters, September 2006 (4 pages)  (preprint on [arXiv](#)).
- [1350] [Detrimental non-Markovian errors for surface code memory](#), by John F Kam, Spiro Gicev, Kavan Modi, Angus Southwell, and Muhammad Usman, arXiv, October 2024 (17 pages).
- [1351] [Fault-Tolerant Quantum Computation with Adversarial Errors](#), by Nikolas P. Breuckmann, Louis Golowich, and Umesh Vazirani, arXiv, August 2026 (132 pages).
- [1352] [Dynamics of superconducting qubit relaxation times](#), by M. Carroll, npj Quantum Information, November 2022 (7 pages).
- [1353] [Fast Estimation of Physical Error Contributions of Quantum Gates](#), by Miha Papič, Adrian Auer, and Inés de Vega, arXiv, July 2023 (6 pages).
- [1354] [The XYZ² hexagonal stabilizer code](#), by Basudha Srivastava et al., Quantum Journal, April 2022 (15 pages).
- [1355] [Sources of decoherence](#), by ETH Zurich, ETH Zurich, 2005 (23 slides).
- [1356] [Impact of ionizing radiation on superconducting qubit coherence](#), by Antti Vepsäläinen, William D. Oliver, et al., arXiv, August 2020 (16 pages).
- [1357] [Crosstalk In Contemporary Quantum Devices](#), by Spiro Gicev, Ben Harper, Haiyue Kang, Muhammad Usman, and Martin Sevir, arXiv, May 2026 (39 pages).

- [1358] [Material Defects in Superconducting Quantum Bits: Origins and Remedies](#), by Jürgen Lisenfeld, Les Houches, 2023 (35 slides).
- [1359] [Developing a Chemical and Structural Understanding of the Surface Oxide in a Niobium Superconducting](#), by Akshay A. Murthy et al., ACS Nano, September 2022 (23 pages)  (preprint on [osti.gov](https://arxiv.org/abs/2209.00000)).
- [1360] [Cosmic-ray-induced correlated errors in superconducting qubit array](#), by Xuegang Li, Hai-Feng Yu, et al., arXiv, June 2025 (9 pages).
- [1361] [Synchronous Detection of Cosmic Rays and Correlated Errors in Superconducting Qubit Arrays](#), by Patrick M. Harrington, William D. Oliver, Joseph A. Formaggio, et al., arXiv, February 2024 (34 pages).
- [1362] [Accelerating Surface Radiation Content to Investigate the Impact of Radon Progeny on Superconducting Qubits](#), by Sagar S. Poudel, Lauren Hsu, et al., arXiv, May 2026 (18 pages).
- [1363] [Iterative Decoding of Stabilizer Codes under Radiation-Induced Correlated Noise](#), by Anuj K. Nayak, Paul G. Baity, Peter J. Love, Nicholas Jeon, Byung-Jun Yoon, Adolfo Hoisie, and Lav R. Varshney, arXiv, March 2026 (14 pages).
- [1364] [Spectroscopic Measurements and Models of Energy Deposition in the Substrate of Quantum Circuits by Natural Ionizing Radiation](#), by Joseph W. Fowler et al., PRX Quantum, November 2024 (21 pages).
- [1365] [High-Temporal-Resolution Measurements of the Impacts of Ionizing Radiation on Superconducting Qubits](#), by Jihee Yang, Thomas J. Carroll, Philip Mason, Robert Schwartz, Kenneth M. O'Hara, Jennifer Lund, Michael Gottschalk, Timothy Stephenson, Lawrence H. Friedman, Francisco Yumiceva, Justin Hackley, Aurelius L. Graninger, Chris Rotella, Pat Warner, Jonathan M. Cochran, Adam V. Bruce, Melody Wagner, James Wenner, Stan Steers, Christopher Moore, Alex Marakov, and Bradley G. Christensen, arXiv, February 2026 (10 pages).
- [1366] [Qubit error bursts in superconducting quantum processors of Quantum Inspire: quasiparticle pumping and anomalous time dependence](#), by G. R. Di Carlo, M. Samiotis, A. Kamalapur, M. Finkel, N. Muthusubramanian, M. W. Beekman, N. Haider, M. S. Moreira, J. F. Marques, B. Segers, S. Vallés-Sanclemente, and Leo DiCarlo, arXiv, March 2026 (16 pages).
- [1367] [Superconducting qubit decoherence correlated with detected radiation events](#), by A. R. Castelli, K. M. Beck, L. D. H. Alegria, L. A. Martinez, K. R. Chaves, S. R. O'Kelley, N. Materise, J. L DuBois, and Y. J. Rosen, arXiv, December 2025 (10 pages).
- [1368] [A Cryogenic Muon Tagging System Based on Kinetic Inductance Detectors for Superconducting Quantum Processors](#), by Ambra Mariani, David van Zanten, et al., arXiv, March 2026 (9 pages).
- [1369] [On-Demand Correlated Errors in Superconducting Qubits from a Particle Accelerator](#), by Thomas McJunkin, A.W. Hunt, Yenuel Jones-Alberty, T.M. Haard, M.K. Spear, James Shackford, Tom Gilliss, Mayra Amezcua, C.A. Watson, T.M. Sweeney, J.A. Hoffmann, and Kevin Schultz, arXiv, March 2026 (16 pages).
- [1370] [Correlated Charge Noise and Relaxation Errors in Superconducting Qubits](#), by C. D. Wilen and Robert McDermott, arXiv, December 2020 (19 pages).
- [1371] [Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits](#), by Matt McEwen, Rami Barends, et al., arXiv, April 2021 (13 pages).
- [1372] [TLS Dynamics in a Superconducting Qubit Due to Background Ionizing Radiation](#), by Ted Thorbeck, Malcolm Carroll, et al., arXiv, October 2022 (14 pages).
- [1373] [Resilience of the surface code to error bursts](#), by Shi Jie Samuel Tan, Christopher A. Pattison, Matt McEwen, and John Preskill, arXiv, June 2024 (13 pages).
- [1374] [Abatement of Ionizing Radiation for Superconducting Quantum Devices](#), by B. Loer, William D. Oliver, M. Warner, et al., arXiv, March 2024 (36 pages).
- [1375] [First Measurement of Correlated Charge Noise in Superconducting Qubits at an Underground Facility](#), by G. Bratrud, D. Bowring, et al., arXiv, June 2024 (12 pages).
- [1376] [Evaluating radiation impact on transmon qubits in above and underground facilities](#), by Francesco De Dominicis, Anna Grassellino, Laura Cardani, et al., arXiv, July 2025 (7 pages).
- [1377] [Reducing the impact of radioactivity on quantum circuits in a deep-underground facility](#), by Laura Cardani, Ioan M. Pop, et al., arXiv, May 2020 (15 pages).
- [1378] [Fermilab goes deep to silence noisy radiation affecting qubits](#), by Hannah H. Means, Physics Today, August 2024.
- [1379] [Phonon downconversion to suppress correlated errors in superconducting qubits](#), by V. Iaia, B. L. T. Plourde, et al., arXiv, November 2022 (24 pages).
- [1380] [Phonon traps reduce the quasiparticle density in superconducting circuits](#), by Fabio Henriques, Ioan M. Pop, et al., arXiv, August 2019 (14 pages).
- [1381] [Fight or Flight: Cosmic Ray-Induced Phonons and the Quantum Surface Code](#), by Bernard Ousmane Sane, Rodney Van Meter, and Michal Hajdušek, arXiv, July 2023 (11 pages).
- [1382] [Distributed quantum error correction for chip-level catastrophic errors](#), by Qian Xu, Andrew N. Cleland, Liang Jiang, et al., arXiv, March 2022 (11 pages).
- [1383] [Resisting high-energy impact events through gap engineering in superconducting qubit arrays](#), by Matt McEwen, Michel Devoret, Alex Opremcak, et al., Physical Review Letters, October 2024.
- [1384] [Recovery dynamics of a gap-engineered transmon after a quasiparticle burst](#), by Heekun Nho, Michel H. Devoret, et al., arXiv, June 2025 (27 pages).
- [1385] [Correlated Error Bursts in a Gap-Engineered Superconducting Qubit Array](#), by Vladislav D. Kurilovich, Alex Opremcak, et al., arXiv, June 2025 (26 pages).
- [1386] [Mitigating cosmic ray-like correlated events with a modular quantum processor](#), by Xuntao Wu, Andrew N. Cleland, et al., arXiv, May 2025 (13 pages).
- [1387] [Intrinsic Multi-Mode Interference for Passive Suppression of Purcell Decay in Superconducting Circuits](#), by Mustafa Bakr, Peter Leek, et al., arXiv, July 2025 (19 pages).
- [1388] [Averting multi-qubit burst errors in surface code magic state factories](#), by Jason D. Chadwick, Frederic T. Chong, et al., arXiv, April 2024 (13 pages).
- [1389] [Q3DE: A fault-tolerant quantum computer architecture for multi-bit burst errors by cosmic rays](#), by Yasunari Suzuki, Teruo Tanimoto, et al., arXiv, December 2024 (16 pages).
- [1390] [Distinguishing types of correlated errors in superconducting qubits](#), by Hannah P. Binney, William D. Oliver, et al., arXiv, March 2026 (14 pages).
- [1391] [Measuring Decoherence Due to Quantum Vacuum Fluctuations](#), by Anirudh Gundhi and Hendrik Ulbricht, arXiv, June 2025 (8 pages).
- [1392] [Observation of quantum many-body effects due to zero point fluctuations in superconducting circuits](#), by Sebastien Leger, Nicolas Roch, et al., arXiv, October 2019 (10 pages).
- [1393] [A model of quantum collapse induced by gravity](#), by F. Laloë, arXiv, February 2020 (14 pages).
- [1394] [Quantum Sensing from Gravity as Universal Dephasing Channel for Qubits](#), by Alexander V. Balatsky, Patrick J. Wong, et al., arXiv, March 2025 (10 pages).
- [1395] [Awesome Quantum Computing Experiments: Benchmarking Experimental Progress Towards Fault-Tolerant Quantum Computation](#), by François-Marie Le Régent, arXiv, July 2025 (29 pages).

- [1396] [An introduction to quantum error correction](#), by Mazyar Mirrahimi, 2018 (31 slides).
- [1397] [Introduction to quantum computing](#), by Anthony Leverrier and Mazyar Mirrahimi, March 2020 (69 slides).
- [1398] [Experimental Quantum Error Correction below the Surface Code Threshold via All-Microwave Leakage Suppression](#), by Tan He et al., Physical Review Letters, December 2025 (7 pages) .
- [1399] [Quantum advantage with shallow circuits](#), by Robert König, 2018 (97 slides).
- [1400] [Helios: A 98-qubit trapped-ion quantum computer with all-to-all connectivity](#), by Anthony Ransford, M.S. Allman, et al., Nature, June 2026 (16 pages).
- [1401] [Efficient measurement of quantum gate error by interleaved randomized benchmarking](#), by Easwar Magesan, Jay M. Gambetta, M. Steffen, et al., arXiv, March 2014 (5 pages).
- [1402] [Quantum Computing: Progress and Prospects](#), by Frank K. Wilhelm et al., BSI, 2018.
- [1403] [Scalable Noise Estimation with Random Unitary Operators](#), by Joseph Emerson, Robert Alicki, and Karol Zyczkowski, arXiv, December 2005 (8 pages).
- [1404] [Randomized benchmarking for individual quantum gates](#), by E. Onorati, A. H. Werner, and J. Eisert, arXiv, December 2018 (13 pages).
- [1405] [A Practical Introduction to Benchmarking and Characterization of Quantum Computers](#), by Akel Hashim, Robin Blume-Kohout, Irfan Siddiqi, et al., arXiv, August 2024 (125 pages).
- [1406] [Efficient learning of quantum noise](#), by Robin Harper, Steven T. Flammia, and Joel J. Wallman, arXiv, April 2021 (17 pages).
- [1407] [Characterization, certification and validation of quantum systems](#), by Martin Kliesch, April 2020 (87 pages).
- [1408] [Characterization of quantum devices](#), by Steve Flammia, University of Sydney, 2017 (118 slides).
- [1409] [Down with T2, long live T3!](#), by Maciej Malinowski, Reading the quantum, September 2025.
- [1410] [Scalability of quantum error mitigation techniques: from utility to advantage](#), by Sergey N. Filippov, Sabrina Maniscalco, and Guillermo García-Pérez, arXiv, March 2024 (36 pages).
- [1411] [On the Importance of Error Mitigation for Quantum Computation](#), by Dorit Aharonov, Maor Shutman, et al., arXiv, March 2025 (28 pages).
- [1412] [Hybrid quantum-classical algorithms and quantum error mitigation](#), by Suguru Endo, Zhenyu Cai, Simon C. Benjamin, and Xiao Yuan, arXiv, November 2020 (39 pages).
- [1413] [Quantum Error Mitigation](#), by Zhenyu Cai, Ryan Babbush, Simon C. Benjamin, Suguru Endo, William J. Huggins, Ying Li, Jarrod R. McClean, and Thomas E. O'Brien, arXiv, December 2023 (43 pages).
- [1414] [Testing platform-independent quantum error mitigation on noisy quantum computers](#), by Vincent Russo, Andrea Mari, Nathan Shammah, Ryan LaRose, and William J. Zeng, arXiv, December 2022 (18 pages).
- [1415] [Simulation and Benchmarking of Real Quantum Hardware](#), by T. Piskor, M. Schöndorf, M. Bauer, D. Smith, T. Ayril, S. Pogorzalek, A. Auer, and M. Papič, arXiv, August 2025 (13 pages).
- [1416] [Fundamental limits of quantum error mitigation](#), by Ryuji Takagi, Mile Gu, et al., npj Quantum Information, September 2022.
- [1417] [What's the difference between error suppression, error mitigation, and error correction?](#), by Matthias Steffen, IBM Research, October 2022.
- [1418] [Error mitigation for short-depth quantum circuits](#), by Kristan Temme, Sergey Bravyi, and Jay M. Gambetta, arXiv, November 2017 (15 pages).
- [1419] [Scalable error mitigation for noisy quantum circuits produces competitive expectation values](#), by Youngseok Kim, Abhinav Kandala, et al., arXiv, August 2021 (7 pages).
- [1420] [Direct Analysis of Zero-Noise Extrapolation: Polynomial Methods, Error Bounds, and Simultaneous Physical-Algorithmic Error Mitigation](#), by Pegah Mohammadipour and Xiantao Li, arXiv, March 2025 (31 pages).
- [1421] [Efficient Variational Quantum Simulator Incorporating Active Error Minimization](#), by Ying Li et al., 2017 (14 pages).
- [1422] [Folding-Free ZNE: A Comprehensive Quantum Zero-Noise Extrapolation Approach for Mitigating Depolarizing and Decoherence Noise](#), by Hrushikesh Patil, Peiyi Li, Ji Liu, and Huiyang Zhou, arXiv, May 2023 (12 pages).
- [1423] [Development of Zero-Noise Extrapolated Projection Based Quantum Algorithm for Accurate Evaluation of Molecular Energetics in Noisy Quantum Devices](#), by Chinmay Shrikhande, Sonaldeep Halder, and Rahul Maitra, arXiv, June 2023 (12 pages).
- [1424] [Error mitigated variational algorithm on a photonic processor](#), by O.V. Borzenkova, S.S. Straupe, et al., arXiv, August 2024 (5 pages).
- [1425] [Energy-Scaled Zero-Noise Extrapolation for Gottesman-Kitaev-Preskill Code](#), by Gui-Zhong Luo and Matthew Otten, arXiv, December 2025 (15 pages).
- [1426] [Demonstrating quantum error mitigation on logical qubits](#), by Aosai Zhang, Zhen Wang, H Wang, et al., Nature Communications, December 2025 (7 pages).
- [1427] [Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors](#), by Ewout van den Berg, Zlatko K. Minev, Abhinav Kandala, and Kristan Temme, arXiv, June 2022 (29 pages).
- [1428] [Unfolding Quantum Computer Readout Noise](#), by Benjamin Nachman, Miroslav Urbanek, Wibe A. de Jong, and Christian W. Bauer, arXiv, May 2020 (13 pages).
- [1429] [Compressed quantum error mitigation](#), by Maurits S. J. Tepaske and David J. Luitz, arXiv, November 2023 (11 pages).
- [1430] [Zeno-Enhanced Probabilistic Error Cancellation with Quantum Error Detection Codes](#), by Yi Yuan, Yuanchen Zhao, and Dong E. Liu, arXiv, May 2026 (33 pages).
- [1431] [CV4Quantum: Reducing the Sampling Overhead in Probabilistic Error Cancellation Using Control Variates](#), by Prasanth Shyamsundar and Wern Yeen Yeong, arXiv, February 2025 (38 pages).
- [1432] [Probabilistic error cancellation for dynamic quantum circuits](#), by Riddhi S. Gupta, Abhinav Kandala, et al., arXiv, December 2023 (16 pages).
- [1433] [Scalable Measurement Error Mitigation via Iterative Bayesian Unfolding](#), by Siddarth Srinivasan, Byron Boots, et al., arXiv, October 2022 (13 pages).
- [1434] [Improving probabilistic error cancellation in the presence of non-stationary noise](#), by Samudra Dasgupta and Travis S. Humble, arXiv, July 2024 (15 pages).
- [1435] [Low bit-flip rate probabilistic error cancellation](#), by Mathys Rennela and Harold Ollivier, arXiv, June 2025 (26 pages).
- [1436] [Theory of quantum error mitigation for non-Clifford gates](#), by David Layden, Bradley Mitchell, and Karthik Siva, arXiv, May 2025 (37 pages).
- [1437] [Co-Designing Error Mitigation and Error Detection for Logical Qubits](#), by Rohan S. Kumar, Takahiro Tsunoda, Sophia H. Xue, Dantong Li, Robert J. Schoelkopf, and Yongshan Ding, arXiv, April 2026 (15 pages).
- [1438] [Machine Learning for Practical Quantum Error Mitigation](#), by Haoran Liao, Zlatko K. Minev, et al., arXiv, November 2024 (11 pages).
- [1439] [Neural networks enable learning of error correction strategies for quantum computers](#), by Thomas Fösel et al., PRX, October 2018.

- [1440] [Error mitigation with Clifford quantum-circuit data](#), by Piotr Czarnik, Lukasz Cincio, et al., arXiv, November 2021 (16 pages).
- [1441] [Improving the efficiency of learning-based error mitigation](#), by Piotr Czarnik, Lukasz Cincio, et al., arXiv, April 2025 (22 pages).
- [1442] [Error mitigation for universal gates on encoded qubits](#), by Christophe Piveteau, Sergey Bravyi, Jay M. Gambetta, Kristan Temme, et al., arXiv, October 2021 (11 pages).
- [1443] [Physics-inspired Machine Learning for Quantum Error Mitigation](#), by Xiao-Yue Xu, Wan-Su Bao, et al., arXiv, January 2025 (25 pages).
- [1444] [Scalable Quantum Error Mitigation with Physically Informed Graph Neural Networks](#), by Huaxin Wang, Xinge Wu, Jiajun Liu, Ruiqing He, Jiandong Shang, Hengliang Guo, and Qiang Chen, arXiv, April 2026 (23 pages).
- [1445] [Deep Learning Approaches to Quantum Error Mitigation](#), by Leonardo Placidi, Ifan Williams, Enrico Rinaldi, Daniel Mills, Cristina Cîrstoiu, Vanya Eccles, and Ross Duncan, arXiv, January 2026 (48 pages).
- [1446] [MCMit: Mid-Circuit Measurement Error Mitigation](#), by Emmanouil Giortamis, Felix Gust, Aleksandra Świerkowska, Sandra Stankovic, Innocenzo Fulginiti, Yanbin Chen, Xiaorang Guo, Benjamin Lienhard, Martin Schulz, and Pramod Bhatotia, arXiv, June 2026 (16 pages).
- [1447] [Scalable tensor-network error mitigation for near-term quantum computing](#), by Sergei Filippov, Guillermo García-Pérez, et al., arXiv, July 2023 (31 pages).
- [1448] [Harnessing Tensor Networks for Error Mitigation in Quantum Computing](#), by Daniel, Algorithmiq, September 2024.
- [1449] [QuantumNAS: Noise-Adaptive Search for Robust Quantum Circuits](#), by Hanrui Wang, Song Han, et al., arXiv, January 2022 (19 pages).
- [1450] [Deep Reinforcement Learning for Quantum State Preparation with Weak Nonlinear Measurements](#), by Riccardo Porotti, Antoine Essig, Benjamin Huard, and Florian Marquardt, arXiv, June 2022 (11 pages).
- [1451] [Reinforcement learning pulses for transmon qubit entangling gates](#), by Ho Nam Nguyen, Markus Schmitt, et al., arXiv, June 2024 (8 pages).
- [1452] [Q-fid: Quantum Circuit Fidelity Improvement with LSTM Networks](#), by Yikai Mao, Shaswot Shresthamali, and Masaaki Kondo, arXiv, May 2025 (13 pages).
- [1453] [Error estimation in current noisy quantum computers](#), by Unai Aseguinolaza, Juan Borge, et al., arXiv, May 2024 (10 pages).
- [1454] [Development and Demonstration of an Efficient Readout Error Mitigation Technique for use in NISQ Algorithms](#), by Andrew Arrasmith, Marco Paini, et al., arXiv, April 2023 (19 pages).
- [1455] [Readout Error Mitigation for Mid-Circuit Measurements and Feedforward](#), by Jin Ming Koh, Dax Enshan Koh, and Jayne Thompson, arXiv, February 2025 (16 pages).
- [1456] [Enhancing Qubit Readout with Autoencoders](#), by Piero Luchi, Francesco Pederiva, et al., arXiv, September 2023 (16 pages).
- [1457] [VarSaw: Application-tailored Measurement Error Mitigation for Variational Quantum Algorithms](#), by Siddharth Dangwal, Frederic T. Chong, et al., arXiv, February 2024 (16 pages).
- [1458] [Error Propagation in NISQ Devices for Solving Classical Optimization Problems](#), by Guillermo González-García et al., PRX Quantum, December 2022 (17 pages).
- [1459] [Exponential Error Suppression for Near-Term Quantum Devices](#), by Bálint Koczor, arXiv, September 2021 (34 pages).
- [1460] [Firing up quantum algorithms - boosting performance up to 9,000× with autonomous error suppression](#), by Michael J. Biercuk, March 2022.
- [1461] [Experimental Benchmarking of an Automated Deterministic Error-Suppression Workflow for Quantum Algorithms](#), by Pranav S. Mundada, Physical Review Applied, August 2023 📄.
- [1462] [Near-Term Quantum Computing Techniques: Variational Quantum Algorithms, Error Mitigation, Circuit Compilation, Benchmarking and Classical Simulation](#), by He-Liang Huang, Gui-Lu Long, et al., arXiv, December 2022 (46 pages).
- [1463] [Shadow Distillation: Quantum Error Mitigation with Classical Shadows for Near-Term Quantum Processors](#), by Alireza Seif, Liang Jiang, et al., arXiv, March 2022 (16 pages).
- [1464] [Virtual Distillation for Quantum Error Mitigation](#), by William J. Huggins, PRX, 2021 (25 pages).
- [1465] [Randomized Compiling for Scalable Quantum Computing on a Noisy Superconducting Quantum Processor](#), by Akel Hashim, 2021 (12 pages).
- [1466] [Quantum Error Mitigation via Quantum-Noise-Effect Circuit Groups](#), by Yusuke Hama and Hirofumi Nishi, arXiv, January 2024 (8 pages).
- [1467] [Mitigating depolarizing noise on quantum computers with noise-estimation circuits](#), by Miroslav Urbanek, Wibe A. de Jong, et al., arXiv, March 2021 (7 pages).
- [1468] [Mitigating errors by quantum verification and post-selection](#), by Rawad Mezher, James Mills, and Elham Kashefi, arXiv, May 2022 (15 pages).
- [1469] [Purification-based quantum error mitigation of pair-correlated electron simulations](#), by T. E. O'Brien, G. Anselmetti, et al., arXiv, October 2022 (10 pages).
- [1470] [Virtual Distillation for Quantum Error Mitigation](#), by William J. Huggins, Ryan Babbush, Jarrod R. McClean, et al., arXiv, August 2021 (26 pages).
- [1471] [Single-shot error mitigation by coherent Pauli checks](#), by Ewout van den Berg, Sergey Bravyi, Jay M. Gambetta, Kristan Temme, et al., arXiv, December 2022 (30 pages).
- [1472] [Quantum error mitigation via matrix product operators](#), by Yuchen Guo and Shuo Yang, arXiv, October 2022 (13 pages).
- [1473] [Classical Noise Inversion: A Practical and Optimal framework for Robust Quantum Applications](#), by Dayue Qin, Ying Li, and You Zhou, arXiv, October 2025 (28 pages).
- [1474] [Quantum Error Mitigation as a Universal Error Reduction Technique: Applications from the NISQ to the Fault-Tolerant Quantum Computing Eras](#), by Yasunari Suzuki, PRX Quantum, March 2022.
- [1475] [Zero noise extrapolation on logical qubits by scaling the error correction code distance](#), by Misty A. Wahl, Andrea Mari, Nathan Shammah, William J. Zeng, and Gokul Subramanian Ravi, arXiv, July 2023 (10 pages).
- [1476] [Mitigating errors in logical qubits](#), by Samuel C. Smith, Benjamin J. Brown, and Stephen D. Bartlett, arXiv, May 2024 (20 pages).
- [1477] [Error Suppression for Arbitrary-Size Black Box Quantum Operations](#), by Gideon Lee, Connor T. Hann, Shruti Puri, S. M. Girvin, and Liang Jiang, arXiv, November 2023 (22 pages).
- [1478] [Unlocking early fault-tolerant quantum computing with mitigated magic dilution](#), by Surabhi Luthra, Earl T. Campbell, et al., arXiv, May 2025 (15 pages).
- [1479] [Infinite Distance Extrapolation: How error mitigation can enhance quantum error correction](#), by George Umbrascu, Oscar Higgott, and Dan E. Browne, arXiv, March 2026 (16 pages).
- [1480] [Multireference error mitigation for quantum computation of chemistry](#), by Hang Zou, Erika Magnusson, Hampus Brunander, Werner Dobrautz, and Martin Rahm, arXiv, May 2025 (16 pages).
- [1481] [Machine Learning Approach towards Quantum Error Mitigation for Accurate Molecular Energetics](#), by Srushti Patil,

- Dibyendu Mondal, and Rahul Maitra, arXiv, April 2025 (14 pages).
- [1482] [Near-Term Fermionic Simulation with Subspace Noise Tailored Quantum Error Mitigation](#), by Miha Papič, Alessio Calzona, et al., arXiv, March 2025 (18 pages).
- [1483] [Prospects of Quantum Error Mitigation for Quantum Signal Processing](#), by Ugnė Liaubaitė and S. E. Skelton, arXiv, May 2025 (8 pages).
- [1484] [Ensemble-Based Quantum Signal Processing for Error Mitigation](#), by Suying Liu, Yulong Dong, Dong An, and Murphy Yuezhen Niu, arXiv, January 2026 (14 pages).
- [1485] [Quantum Error Mitigation for Sampling Algorithms](#), by Kecheng Liu and Zhenyu Cai, arXiv, February 2025 (13 pages).
- [1486] [Quantum error mitigation in quantum annealing](#), by Mohammad H. Amin, Marcel Franz, et al., arXiv, November 2023 (10 pages).
- [1487] [Quantum error mitigation in quantum annealing](#), by Jack Raymond, Jacek Dziarmaga, et al., npj Quantum Information, March 2025.
- [1488] [Error mitigation of shot-to-shot fluctuations in analog quantum simulators](#), by Thomas Steckmann, Alexander Schuckert, et al., arXiv, June 2025 (20 pages).
- [1489] [Extending the computational reach of a superconducting qutrit processor](#), by Noah Goss, Irfan Siddiqi, et al., npj Quantum Information, October 2024.
- [1490] [Benchmarking Noisy Intermediate Scale Quantum Error Mitigation Strategies for Ground State Preparation of the HCl Molecule](#), by Tim Weaving, Peter V. Coveney, et al., arXiv, March 2023 (18 pages).
- [1491] [Hypothesis Testing for Error Mitigation: How to Evaluate Error Mitigation](#), by Abdullah Ash Saki, George Umbraescu, et al., arXiv, January 2023 (19 pages).
- [1492] [Noise tailoring for error mitigation and for diagnosing digital quantum computers](#), by Thibault Scoquart, Hugo Perrin, and Kyrylo Snizhko, arXiv, March 2026 (27 pages).
- [1493] [Making quantum error mitigation practical](#), by Will Zeng et al., Unitary Fund, May 2023.
- [1494] [Error statistics and scalability of quantum error mitigation formulas](#), by Dayue Qin, Yanzhu Chen, and Ying Li, arXiv, April 2023 (25 pages).
- [1495] [Exponentially tighter bounds on limitations of quantum error mitigation](#), by Yihui Quek et al., Nature, October 2022 (23 pages).
- [1496] [Error Mitigation Thresholds in Noisy Random Quantum Circuits](#), by Pradeep Niroula, Sarang Gopalakrishnan, and Michael J. Gullans, arXiv, June 2024 (11 pages).
- [1497] [The Exponential Deviation Induced by Quantum Readout Error Mitigation](#), by Yibin Guo, Hai-Feng Yu, et al., arXiv, October 2025 (10 pages).
- [1498] [Limitations of probabilistic error cancellation for open dynamics beyond sampling overhead](#), by Yue Ma and M. S. Kim, arXiv, January 2024 (13 pages).
- [1499] [Can Error Mitigation Improve Trainability of Noisy Variational Quantum Algorithms?](#), by Samson Wang, Marco Cerezo, Patrick J. Coles, et al., arXiv, March 2024 (29 pages).
- [1500] [A polynomial-time classical algorithm for noisy quantum circuits](#), by Thomas Schuster, Chao Yin, Xun Gao, and Norman Y. Yao, arXiv, October 2024 (11 pages).
- [1501] [Purification-based quantum error mitigation of pair-correlated electron simulations](#), by T. E. O'Brien, Ryan Babbush, N. C. Rubin, et al., Nature Physics, October 2023 (preprint on doi.org).
- [1502] [Error mitigation and circuit division for early fault-tolerant quantum phase estimation](#), by Alicja Dutkiewicz, Thomas E. O'Brien, et al., arXiv, February 2025 (37 pages).
- [1503] [Quantum Error Suppression with Subgroup Stabilisation](#), by Bo Yang, Elham Kashefi, Dominik Leichtle, and Harold Ollivier, arXiv, June 2024 (32 pages).
- [1504] [Resource-efficient context-aware dynamical decoupling embedding for arbitrary large-scale quantum algorithms](#), by Paul Coote, Roman Dimov, Smarak Maity, Gavin S. Hartnett, Michael J. Biercuk, and Yuval Baum, arXiv, September 2024 (12 pages).
- [1505] [Halving the Cost of Quantum Algorithms with Randomization](#), by John M. Martyn and Patrick Rall, arXiv, March 2025 (19 pages).
- [1506] [Improving error suppression with noise-aware decoding](#), by Evan T. Hockings, Andrew C. Doherty, and Robin Harper, arXiv, April 2025 (5 pages).
- [1507] [Impact of time-retarded noise on dynamical decoupling schemes for qubits](#), by Kiyoto Nakamura and Joachim Ankerhold, arXiv, February 2025 (15 pages).
- [1508] [Effects of Dynamical Decoupling and Pulse-level Optimizations on IBM Quantum Computers](#), by Siyuan Niu and Aida Todri-Sanial, arXiv, September 2022 (14 pages).
- [1509] [Analyzing Strategies for Dynamical Decoupling Insertion on IBM Quantum Computer](#), by Siyuan Niu and Aida Todri-Sanial, arXiv, April 2022 (6 pages).
- [1510] [Suppressing errors with dynamical decoupling using pulse control on Amazon Braket](#), by Palash Goiporia et al., PRA, December 2022.
- [1511] [Demonstration of algorithmic quantum speedup](#), by Bibek Pokharel and Daniel A. Lidar, arXiv, July 2022 (12 pages).
- [1512] [Designing dynamically corrected gates robust to multiple noise sources using geometric space curves](#), by Hunter T. Nelson, Edwin Barnes, et al., arXiv, November 2022 (12 pages).
- [1513] [Dynamically corrected gates from geometric space curves](#), by Edwin Barnes, Fei Zhuang, et al., arXiv, March 2021 (25 pages).
- [1514] [Study and Application of the Space Curve Quantum Control Formalism](#), by Fei Zhuang, May 2023 (205 pages).
- [1515] [Quantum Error Management in Practice: A Cross-Stack Benchmark](#), by Daniel Sierra-Sosa, Begonya Garcia-Zapirain, Cristian Marquez, and Kelly Garces, arXiv, August 2026 (21 pages).
- [1516] [Quantum optimization using a 127-qubit gate-model IBM quantum computer can outperform quantum annealers for nontrivial binary optimization problems](#), by Natasha Sachdeva, Michael J. Biercuk, Yuval Baum, et al., arXiv, October 2024 (19 pages).
- [1517] [A comment on comparing optimization on D-Wave and IBM quantum processors](#), by Catherine C. McGeoch, Kevin Chern, Pau Farré, and Andrew K. King, arXiv, June 2024 (5 pages).
- [1518] [Protecting Expressive Circuits with a Quantum Error Detection Code](#), by Chris N. Self, Marcello Benedetti, and David Amaro, arXiv, July 2024 (16 pages).
- [1519] [Opportunities and challenges in scaling quantum error detection on hardware](#), by Yanis Le Fur, Ethan Egger, Hong-Ye Hu, Vincent Russo, William J. Zeng, and Ryan LaRose, arXiv, May 2026 (11 pages).
- [1520] [Scalable Postselection of Quantum Resources](#), by J. Wilson Staples, Winston Fu, and Jeff D. Thompson, arXiv, March 2026 (9 pages).
- [1521] [Toward Fault-Tolerant Variational Optimization: QAOA under \$\[\[4,2,2\]\]\$ Error Detection](#), by Matteo Robert Child, Emanuele Dri, Giacomo Vitali, Chiara Vercellino, and Alberto Leporati, arXiv, September 2026 (12 pages).
- [1522] [Realizing Error Suppression in Partially Fault-Tolerant Quantum Simulations with IBM Quantum Computers](#), by Henry Froland, Dorota M. Grabowska, Sebastian Grieninger, Jeremy Hartse, Anne L. Lashbrook, Zhiyao Li, Ziyuan Li, Sarah J. M. Powell, Martin J. Savage, Xiao

- jun Yao, and Nikita A. Zemlevskiy, arXiv, July 2026 (42 pages).
- [1523] [Practical Error Suppression and Mitigation for Reliable Quantum Computing](#), by Han-Ze Li, Mengjie Yang, Xianquan Yan, Dax Enshan Koh, Ching Hua Lee, and Ruizhe Shen, arXiv, August 2026 (122 pages).
 - [1524] [A Tutorial on Quantum Error Correction](#), by Andrew M. Steane, 2006 (24 pages).
 - [1525] [Quantum Error Correction For Dummies](#), by Avimita Chatterjee, Koustubh Phalak, and Swaroop Ghosh, arXiv, April 2023 (12 pages).
 - [1526] [Quantum Circuits for Stabilizer Error Correcting Codes: A Tutorial](#), by Arijit Mondal and Keshab K. Parhi, arXiv, September 2023 (15 pages).
 - [1527] [Decoding algorithms for surface codes](#), by Antonio deMartioli, Patricio Fuentes, Román Orús, Pedro M. Crespo, and Josu Etxezarreta Martinez, arXiv, September 2024 (59 pages).
 - [1528] [Impacts of Decoder Latency on Utility-Scale Quantum Computer Architectures](#), by Abdullah Khalid, Yonatan Cohen, Pooya Ronagh, et al., arXiv, November 2025 (23 pages).
 - [1529] [To reset, or not to reset – that is the question](#), by György P. Gehér, Marcin Jastrzebski, Earl T. Campbell, and Ophelia Crawford, arXiv, March 2025 (13 pages).
 - [1530] [A High Performance Compiler for Very Large Scale Surface Code Computations](#), by George Watkins, Alexandru Paler, et al., arXiv, May 2024 (16 pages).
 - [1531] [Dependency-Aware Compilation for Surface Code Quantum Architectures](#), by Abtin Molavi, Amanda Xu, Swamit Tannu, and Aws Albarghouti, arXiv, April 2025 (38 pages).
 - [1532] [A Theory of Quantum Error-Correcting Codes](#), by Emanuel Knill and Raymond Laflamme, arXiv, April 1996 (34 pages).
 - [1533] [Mixed State Entanglement and Quantum Error Correction](#), by Charles H. Bennett, David P. DiVincenzo, John A. Smolin, and William K. Wootters, arXiv, August 1996 (82 pages).
 - [1534] [Quantum Information Processing and Quantum Error Correction. An Engineering Approach](#), by Ivan Djordjevic, April 2021 (575 pages).
 - [1535] [Modular fault-tolerant quantum computing on a non-CSS code](#), by Robert Freund, Friederike Butt, César Benito, Ivan Pogorelov, Marcel Meyer, Alex Steiner, Alejandro Bermudez, Markus Müller, and Thomas Monz, arXiv, September 2026 (24 pages).
 - [1536] [Classical product code constructions for quantum Calderbank-Shor-Steane codes](#), by Dimitar Ostrev, Balazs Matuz, et al., arXiv, July 2024 (27 pages).
 - [1537] [Engineering Quantum Error Correction Codes Using Evolutionary Algorithms](#), by Mark Webster and Dan Browne, arXiv, September 2024 (14 pages).
 - [1538] [Concatenate codes, save qubits](#), by Satoshi Yoshida, Shiro Tamiya, and Hayata Yamasaki, npj Quantum Information, May 2025 (9 pages).
 - [1539] [Quantum Error Correction for Beginners](#), by Simon J. Devitt, Kae Nemoto, and William J. Munro, arXiv, June 2013 (35 pages).
 - [1540] [Quantum error correction \(QEC\)](#), by Alexander Korotkov, 2017 (39 slides).
 - [1541] [Quantum Error Detection Without Using Ancilla Qubits](#), by Nicolas J. Guerrero and David E. Weeks, arXiv, April 2022 (8 pages).
 - [1542] [Fault-tolerant measurement-free quantum error correction with multiqubit gates](#), by Michael A. Perlin, Mark Saffman, et al., PRA, December 2023 (14 pages) (preprint on arXiv).
 - [1543] [Measurement-free, scalable and fault-tolerant universal quantum computing](#), by Friederike Butt, Markus Müller, et al., arXiv, October 2024 (16 pages).
 - [1544] [Measurement-free fault-tolerant quantum error correction in near-term devices](#), by Sascha Heußen, David F. Locher, and Markus Müller, PRX Quantum, February 2024 (25 pages).
 - [1545] [Demonstration of measurement-free universal fault-tolerant quantum computation](#), by Friederike Butt, Thomas Monz, Markus Müller, et al., arXiv, June 2025 (19 pages).
 - [1546] [Protecting a Bosonic Qubit with Autonomous Quantum Error Correction](#), by Jeffrey M. Gertler, Jens Koch, Chen Wang, et al., arXiv, October 2020 (5 pages).
 - [1547] [Autonomous quantum error correction and quantum computation](#), by José Lebreuilly, Steven M. Girvin, Liang Jiang, et al., arXiv, March 2021 (18 pages).
 - [1548] [Autonomous quantum error correction with superconducting qubits](#), by Joachim Cohen, 2017 (164 pages).
 - [1549] [Surface codes: Towards practical large-scale quantum computation](#), by Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland, Physical Review A, September 2012 (54 pages) (preprint on arXiv).
 - [1550] [Quantum codes on a lattice with boundary](#), by S. B. Bravyi and A. Yu. Kitaev, arXiv, November 1998 (6 pages).
 - [1551] [Scalable Quantum Error Correction for Surface Codes using FPGA](#), by Namitha Liyanage, Lin Zhong, et al., arXiv, May 2023 (12 pages).
 - [1552] [On Error Thresholds for Pauli Channels: Some answers with many more questions](#), by Avantika Agarwal, Alan Bu, Amolak Ratan Kalra, Debbie Leung, Luke Schaeffer, and Graeme Smith, arXiv, March 2026 (34 pages).
 - [1553] [Approaches to Quantum Error Correction](#), by Julia Kempe, arXiv, December 2006 (29 pages).
 - [1554] [Quantum Error Correction for Quantum Memories](#), by Barbara M. Terhal, arXiv, April 2015 (47 pages).
 - [1555] [Introduction to Quantum Error Correction and Fault Tolerance](#), by Steven M. Girvin, arXiv, January 2023 (102 pages).
 - [1556] [Surviving as a Quantum Computer in a Classical World](#), by Daniel Gottesman, February 2024 (191 pages).
 - [1557] [Quantum Error Correction for Kids](#), by Richard A. Wolf, arXiv, May 2024 (5 pages).
 - [1558] [Decoding Correlated Errors in Quantum LDPC Codes](#), by Arshpreet Singh Maan, Valentin Savin, et al., Nature Communications, March 2026 (13 pages).
 - [1559] [Some Progress on Quantum Error Correction for Discrete and Continuous Error Models](#), by Jincao Li, 2020 (15 pages).
 - [1560] [The Error Correction Zoo](#), by Victor V. Albert and Philippe Faist, 2025.
 - [1561] [Handbook of Error-Correcting Codes](#), by Victor V. Albert and Philippe Faist, arXiv, June 2026 (676 pages).
 - [1562] [Demonstrating experimentally the encoding and dynamics of an error-correctable logical qubit on a hyperfine-coupled nuclear spin qudit](#), by Sumin Lim, Arzhang Ardavan, et al., arXiv, January 2025 (41 pages).
 - [1563] [Weight Reduced Stabilizer Codes with Lower Overhead](#), by Eric Sabo, Guillaume Dauphinais, et al., arXiv, February 2024 (41 pages).
 - [1564] [Universal graph representation of stabilizer codes](#), by Andrey Boris Khesin, Jonathan Z. Lu, and Peter W. Shor, arXiv, February 2025 (73 pages).
 - [1565] [Quantum Low-Density Parity-Check Codes](#), by Bane Vasic, Valentin Savin, Nithin Raveendran, et al., arXiv, October 2025 (40 pages).
 - [1566] [Quantum Low-Density Parity-Check \(qLDPC\) Codes](#), by Marin Ivezic, PostQuantum, March 2026.
 - [1567] [Quantum LDPC codes with design rate 1/5 and good performance below 1000 physical qubits](#), by Yifan Hong, arXiv, July 2026 (69 pages).

- [1568] [Asymptotically Good Quantum and Locally Testable Classical LDPC Codes](#), by Pavel Panteleev and Gleb Kalachev, arXiv, January 2022 (51 pages).
- [1569] [Quantum LDPC codes with positive rate and minimum distance proportional to \$n^{1/2}\$](#) , by Jean-Pierre Tillich, Gilles Zemor, et al., IEEE International Symposium on Information Theory, June 2009.
- [1570] [Fault-Tolerant Quantum Computation with Constant Overhead](#), by Daniel Gottesman, arXiv, July 2014 (32 pages).
- [1571] [The Physics of \(good\) LDPC Codes I. Gauging and dualities](#), by Tibor Rakovszky and Vedika Khemani, arXiv, October 2023 (38 pages).
- [1572] [The Physics of \(good\) LDPC Codes II. Product constructions](#), by Tibor Rakovszky and Vedika Khemani, arXiv, February 2024 (53 pages).
- [1573] [Quantum Low-Density Parity-Check Codes](#), by Nikolas P. Breuckmann and Jens Niklas Eberhardt, arXiv, October 2021 (17 pages).
- [1574] [Opportunities and Challenges in Fault-Tolerant Quantum Computation](#), by Daniel Gottesman, arXiv, October 2022 (24 pages).
- [1575] [Constant-Overhead Quantum Error Correction with Thin Planar Connectivity](#), by Maxime A. Tremblay et al., PRL, July 2022 (7 pages) 📄 (preprint on [arXiv](#)).
- [1576] [Networked Realization of Quantum LDPC Codes](#), by Swayangprabha Shaw and Narayanan Rengaswamy, arXiv, April 2026 (9 pages).
- [1577] [Space-Time Tradeoffs of Pauli-Based Computation in Distributed qLDPC Architectures](#), by Naphan Benchasat-tabuse, Michal Hajdušek, and Rodney Van Meter, arXiv, May 2026 (12 pages).
- [1578] [Topological and Subsystem Codes on Low-Degree Graphs with Flag Qubits](#), by Christopher Chamberland et al., PRX, January 2020 (19 pages).
- [1579] [Codesign of quantum error-correcting codes and modular chiplets in the presence of defects](#), by Sophia Fuhui Lin, Benjamin J. Brown, et al., arXiv, March 2024 (16 pages).
- [1580] [High-threshold and low-overhead fault-tolerant quantum memory](#), by Sergey Bravyi, Theodore J. Yoder, et al., arXiv, February 2024 (36 pages).
- [1581] [Computing with error-corrected quantum computers](#), by Guanyu Zhu et al., IBM, February 2024.
- [1582] [Efficient Routing of Quantum LDPC Codes on Programmable 2D Toric Architectures](#), by Kun Liu, Takahiro Tsunoda, Sophia H. Xue, Evan McKinney, Zeyuan Zhou, Shifan Xu, Robert J. Schoelkopf, and Yongshan Ding, arXiv, April 2026 (13 pages).
- [1583] [High-rate quantum LDPC codes for long-range-connected neutral atom registers](#), by Laura Pecorari, Sven Jandura, Gavin K. Brennen, and Guido Pupillo, Nature Communications, January 2025.
- [1584] [Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays](#), by Qian Xu, J. Pablo Bonilla Ataides, Nithin Raveendran, Bane Vasic, Mikhail D. Lukin, Liang Jiang, Hengyun Zhou, et al., Nature Physics, April 2024 📄 (preprint on [arXiv](#)).
- [1585] [Ambiguity Clustering: an accurate and efficient decoder for qLDPC codes](#), by Stasiu Wolanski and Ben Barber, arXiv, January 2025 (19 pages).
- [1586] [Toward Low-latency Iterative Decoding of QLDPC Codes Under Circuit-Level Noise](#), by Anqi Gong, Sebastian Cammerer, and Joseph M. Renes, arXiv, March 2024 (4 pages).
- [1587] [Belief Propagation Decoding of Quantum LDPC Codes with Guided Decimation](#), by Hanwen Yao, Henry D. Pfister, et al., arXiv, June 2024 (19 pages).
- [1588] [Collective Bit Flipping-Based Decoding of Quantum LDPC Codes](#), by Dimitris Chytas, Nithin Raveendran, and Bane Vasic, arXiv, June 2024 (13 pages).
- [1589] [Bias-tailored single-shot quantum LDPC codes](#), by Shixin Wu, Todd A. Brun, and Daniel A. Lidar, arXiv, July 2025 (20 pages).
- [1590] [Analog Information Decoding of Bosonic Quantum Low-Density Parity-Check Codes](#), by Lucas Berent et al., PRX Quantum, May 2024 (33 pages).
- [1591] [Toward a 2D Local Implementation of Quantum LDPC Codes](#), by Noah Berthussen, Andrew M. Childs, Daniel Gottesman, et al., arXiv, January 2025 (18 pages).
- [1592] [Transversal non-Clifford gates for quantum LDPC codes on sheaves](#), by Ting-Chun Lin, arXiv, October 2024 (34 pages).
- [1593] [Quantum LDPC Codes with Transversal Non-Clifford Gates via Products of Algebraic Codes](#), by Louis Golowich and Ting-Chun Lin, arXiv, October 2024 (63 pages).
- [1594] [Lift-Connected Surface Codes](#), by Josias Old, Manuel Risper, and Markus Müller, arXiv, August 2024 (23 pages).
- [1595] [Mirror codes: High-threshold quantum LDPC codes beyond the CSS regime](#), by Andrey Boris Khesin and Jonathan Z. Lu, arXiv, March 2026 (40 pages).
- [1596] [Self-dual Stacked Quantum Low-Density Parity-Check Codes](#), by Ze-Chuan Liu, Chong-Yuan Xu, and Yong Xu, arXiv, February 2026 (8 pages).
- [1597] [Non-Abelian Quantum Low-Density Parity Check Codes and Non-Clifford Operations from Gauging Logical Gates via Measurements](#), by Maine Christos, Chiu Fan Bowen Lo, Vedika Khemani, and Rahul Sahay, arXiv, February 2026 (36 pages).
- [1598] [Spatially-Coupled QLDPC Codes](#), by Siyi Yang and Robert Calderbank, arXiv, April 2025 (36 pages).
- [1599] [Spacetime codes of Clifford circuits](#), by Nicolas Delfosse and Adam Paetznick, arXiv, May 2023 (33 pages).
- [1600] [Hierarchical memories: Simulating quantum LDPC codes with local gates](#), by Christopher A. Pattison, Anirudh Krishna, and John Preskill, arXiv, April 2025 (70 pages).
- [1601] [Large-Language-Model Discovery of Quantum LDPC Codes through Structured Concept Evolution](#), by Zidu Liu and Florian Marquardt, arXiv, June 2026 (17 pages).
- [1602] [Sparse Graph Codes for Quantum Error-Correction](#), by David J.C. MacKay, Graeme Mitchison, and Paul L. McFadden, arXiv, February 2004 (42 pages).
- [1603] [High-threshold and low-overhead fault-tolerant quantum memory](#), by Sergey Bravyi, Theodore J. Yoder, et al., Nature, March 2024 (11 pages).
- [1604] [Tradeoffs on the volume of fault-tolerant circuits](#), by Anirudh Krishna and Gilles Zémor, arXiv, October 2025 (17 pages).
- [1605] [Accelerating Fault-Tolerant Quantum Computation with Good Quantum Low-Density Parity-Check Codes](#), by Guo Zhang, PRX Quantum, July 2026 (49 pages).
- [1606] [Fast correlated decoding of transversal logical algorithms](#), by Madelyn Cain, Dolev Bluvstein, Chen Zhao, Shouzen Gu, Nishad Maskara, Marcin Kalinowski, Alexandra A. Geim, Aleksander Kubica, Mikhail D. Lukin, and Hengyun Zhou, arXiv, June 2025 (11 pages).
- [1607] [Placing and routing quantum LDPC codes in multilayer superconducting hardware](#), by Melvin Mathews, Lukas Pahl, David Pahl, Vaishnavi L. Addala, Catherine Tang, William D. Oliver, and Jeffrey A. Grover, npj Quantum Information, April 2026 (15 pages).
- [1608] [Routing Codes: High-Rate Quantum LDPC Codes with Short, Parallel Non-Local Connectivity](#), by Jiaxuan Zhang, Zhao-Yun Chen, Peng Duan, Jia-Ning Li, Tian-Hao Wei, Qing-Yang Hou, Wei-Cheng Kong, Yu-Chun Wu, and Guo-Ping Guo, arXiv, June 2026 (14 pages).
- [1609] [Electrical Interconnects for Silicon Spin Qubits](#), by Christopher David White, Anthony Sigillito, and Michael J. Gullans, Physical Review B, February 2026 📄 (preprint on [arXiv](#)).

- [1610] [Barbell Codes: qLDPC Codes for Superconducting Quantum Hardware](#), by Shin Ho Choe, Vincent Steffan, Florian Vigneau, Pedro Parrado-Rodríguez, Hsiang-Sheng Ku, Martin Leib, Francisco Revson Fernandes Pereira, and Fedor Simkovic IV, arXiv, June 2026 (27 pages).
- [1611] [QGPU: Parallel logic in quantum LDPC codes](#), by Boren Gu, Andy Zeyi Liu, Armanda O. Quintavalle, Qian Xu, Jens Eisert, and Joschka Roffe, arXiv, March 2026 (16 pages).
- [1612] [Quantum Tanner codes](#), by Anthony Leverrier and Gilles Zémor, arXiv, September 2022 (35 pages).
- [1613] [Efficient decoding up to a constant fraction of the code length for asymptotically good quantum codes](#), by Anthony Leverrier and Gilles Zémor, arXiv, October 2022 (43 pages).
- [1614] [Single-shot decoding of good quantum LDPC codes](#), by Shouzheng Gu, Aleksander Kubica, et al., arXiv, April 2024 (36 pages).
- [1615] [Small quantum Tanner codes from left–right Cayley complexes](#), by Anthony Leverrier, Wouter Rozendaal, and Gilles Zémor, arXiv, December 2025 (13 pages).
- [1616] [Tour de gross: A modular quantum computer based on bivariate bicycle codes](#), by Theodore J. Yoder, Eddie Schoute, Patrick Rall, Emily Pritchett, Jay M. Gambetta, Andrew W. Cross, Malcolm Carroll, and Michael E. Beverland, arXiv, June 2025 (68 pages).
- [1617] [Lowering Connectivity Requirements For Bivariate Bicycle Codes Using Morphing Circuits](#), by Mackenzie H. Shaw and Barbara M. Terhal, arXiv, March 2025 (29 pages).
- [1618] [Logical Operators and Fold-Transversal Gates of Bivariate Bicycle Codes](#), by Jens Niklas Eberhardt and Vincent Steffan, arXiv, July 2024 (21 pages).
- [1619] [Demonstration of low-overhead quantum error correction codes](#), by Ke Wang, Dong-Ling Deng, et al., arXiv, May 2025 (26 pages).
- [1620] [Anyon Theory and Topological Frustration of High-Efficiency Quantum LDPC Codes](#), by Keyang Chen, Hao Song, et al., arXiv, March 2025 (13 pages).
- [1621] [Compare the Pair: Rotated vs. Unrotated Surface Codes at Equal Logical Error Rates](#), by Anthony Ryan O’Rourke and Simon Devitt, arXiv, October 2024 (12 pages).
- [1622] [Demonstrating dynamic surface codes](#), by Alec Eickbusch et al., arXiv, June 2025 (11 pages).
- [1623] [Relaxing Hardware Requirements for Surface Code Circuits using Time-dynamics](#), by Matt McEwen, Dave Bacon, and Craig Gidney, arXiv, September 2023 (55 pages).
- [1624] [The XZZX Surface Code](#), by J. Pablo Bonilla Ataides, David K. Tuckett, Stephen D. Bartlett, Steven T. Flammia, and Benjamin J. Brown, Nature Communications, April 2021 (12 pages)  (preprint on [arXiv](#)).
- [1625] [How to Build a Quantum Supercomputer: Scaling from Hundreds to Millions of Qubits](#), by Masoud Mohseni, Alan Ho, Raymond G. Beausoleil, John M. Martinis, et al., arXiv, March 2026 (91 pages).
- [1626] In quantum error correction, a code check weight is the number of data qubits checked at once by a single parity-measurement circuit. It indicates how many physical qubits must interact during error detection and determines how easily hardware noise can spread into multiple errors. For a planar rotated surface code, the check weight is 4 (bulk) or 2 (edges) due to nearest neighbor connectivity. Color codes have a checkweight of 6 in hexagonal lattices or a mix of 4 and 8 in square-octagon lattices.
- [1627] [Surface-Code Thresholds and Qubit Footprints in Shuttling-Based Spin-Qubit Railways](#), by Arun John Moncy, Reza Dastbasteh, Josu Etxezarreta Martinez, Ryo Nagai, Pedro M. Crespo, Normann Mertig, Charles Smith, and Ruben M. Otxoa, arXiv, May 2026 (14 pages).
- [1628] [Riverlane Quantum Error Correction Report 2024](#), by Riverlane, October 2024 (67 pages).
- [1629] [An introduction to the surface code](#), by Andrew Cleland, 2022 (68 pages).
- [1630] [Long-range-enhanced surface codes](#), by Yifan Hong, Matteo Marinelli, Adam M. Kaufman, and Andrew Lucas, arXiv, October 2024 (23 pages).
- [1631] [An Architecture for Improved Surface Code Connectivity in Neutral Atoms](#), by Joshua Vizslai, Frederic T. Chong, et al., arXiv, September 2023 (13 pages).
- [1632] [Yoked surface codes](#), by Craig Gidney, Michael Newman, Peter Brooks, and Cody Jones, arXiv, December 2023 (23 pages).
- [1633] [Inplace Access to the Surface Code Y Basis](#), by Craig Gidney, arXiv, April 2024 (25 pages).
- [1634] [The Pinnacle Architecture: Reducing the cost of breaking RSA-2048 to 100 000 physical qubits using quantum LDPC codes](#), by Paul Webster, Lawrence Z. Cohen, et al., arXiv, February 2026 (22 pages).
- [1635] [Securing Elliptic Curve Cryptocurrencies against Quantum Vulnerabilities: Resource Estimates and Mitigations](#), by Ryan Babbush, Adam Zalcman, Craig Gidney, Michael Broughton, Tanuj Khattar, Hartmut Neven, Thiago Bergamaschi, Justin Drake, and Dan Boneh, arXiv, March 2026 (57 pages).
- [1636] Here is the exact doubt expressed by Google in their March 2026 paper: “Similarly, the recently introduced “Pinnacle” architecture [118] suggests breaking 2048-bit RSA with fewer than one hundred thousand superconducting qubits, but relies on even more challenging degree-ten, non-planar connectivity for codes that currently lack extensive simulations characterizing their performance, or efficient decoders, and would require nearly a month of runtime. The physical complexity of devices meeting these more demanding specifications makes them potentially more difficult to fabricate with sufficiently low physical error rates than larger devices with more benign engineering demands. Consequently, while a Pinnacle-style device would theoretically require even fewer than one-hundred thousand superconducting qubits to break 256-bit ECDLP, a planar degree-four device with half a million qubits remains the better understood and also likely more feasible engineering path”.
- [1637] [A Denser Planar Surface Code](#), by Guang Hao Low, William J. Huggins, Dominic W. Berry, Tanuj Khattar, Alec F. White, Nicholas C. Rubin, and Ryan Babbush, arXiv, 2026 (105 pages).
- [1638] [The Pinnacle Architecture with fixed connectivity of degree eight](#), by Paul Webster, Tom Peham, and Lawrence Z. Cohen, arXiv, September 2026 (8 pages).
- [1639] [Teleportation transition of surface codes on a superconducting quantum processor](#), by Yiren Zou, H. Wang, et al., arXiv, February 2026 (20 pages).
- [1640] [Hierarchical Logical Processor on the Rotated Surface Code with Shuttle Buses](#), by Zi-Han Chen, Ming-Cheng Chen, Chao-Yang Lu, and Jian-Wei Pan, arXiv, June 2026 (37 pages).
- [1641] [The XZZX surface code](#), by J. Pablo Bonilla Ataides, David K. Tuckett, Stephen D. Bartlett, Steven T. Flammia, and Benjamin J. Brown, Nature Communications, April 2021  (preprint on [arXiv](#)).
- [1642] [Performance Analysis of a Repetition Cat Code Architecture: Computing 256-bit Elliptic Curve Logarithm in 9 Hours with 126 133 Cat Qubits](#), by Élie Gouzien, Physical Review Letters, July 2023  (preprint on [HAL](#)).
- [1643] [Tailored XZZX codes for biased noise](#), by Qian Xu, Nam Mannucci, Alireza Seif, Aleksander Kubica, Steven T. Flammia, and Liang Jiang, arXiv, March 2022 (16 pages).
- [1644] [Quantum Error Correction with Biased Noise](#), by Peter Brooks, Caltech, 2013 (198 pages).
- [1645] [Quantum error correction below the surface code threshold](#), by Rajeev Acharya, Frank Arute, Michel Devoret, Edward

- Farhi, Craig Gidney, William D. Oliver, Pedram Roushan, et al., *Nature*, December 2024.
- [1646] [The domain wall color code](#), by Konstantin Tiurev, Jan-Michael Reiner, et al., arXiv, February 2024 (12 pages).
- [1647] [Facilitating Practical Fault-tolerant Quantum Computing Based on Color Codes](#), by Jiaxuan Zhang, Yu-Chun Wu, and Guo-Ping Guo, arXiv, June 2024 (20 pages).
- [1648] [The ABCs of the color code](#), by Aleksander Marek Kubica, 2018 (205 pages).
- [1649] [The cost of universality: A comparative study of the overhead of state distillation and code switching with color codes](#), by Michael E. Beverland, Aleksander Kubica, and Krysta M. Svore, arXiv, January 2021 (69 pages).
- [1650] [New circuits and an open source decoder for the color code](#), by Craig Gidney and Cody Jones, arXiv, December 2023 (26 pages).
- [1651] [Improving threshold for fault-tolerant color code quantum computing by flagged weight optimization](#), by Yugo Takada and Keisuke Fujii, arXiv, September 2024 (16 pages).
- [1652] [Domain Wall Color Code](#), by Konstantin Tiurev, Jens Eisert, Markus S. Kesselring, Jan-Michael Reiner, et al., PRL, June 2023 (17 pages)  (preprint on arXiv).
- [1653] [A flow-map model for analyzing pseudothresholds in fault-tolerant quantum computing](#), by Krysta M. Svore, A. W. Cross, Isaac L. Chuang, and A. V. Aho, arXiv, January 2006 (20 pages).
- [1654] [Improving Threshold for Fault-Tolerant Color-Code Quantum Computing by Flagged Weight Optimization](#), by Yugo Takada, PRX Quantum, September 2024.
- [1655] [Very low overhead fault-tolerant magic state preparation using redundant ancilla encoding and flag qubits](#), by Christopher Chamberland and Kyungjoo Noh, npj Quantum Information, October 2020 (12 pages).
- [1656] [Low-overhead quantum computing with the color code](#), by Felix Thomsen, Physical Review Research, November 2024 (12 pages).
- [1657] [Quantum Rainbow Codes](#), by Thomas R. Scruby, Arthur Pesah, and Mark Webster, arXiv, October 2024 (19 pages).
- [1658] [Dynamically Generated Logical Qubits](#), by Matthew B. Hastings and Jeongwan Haah, arXiv, October 2021 (18 pages).
- [1659] [Optimization of the surface code design for Majorana-based qubits](#), by Rui Chao, Michael E. Beverland, Nicolas Delfosse, et al., Quantum Journal, 2020 (19 pages).
- [1660] [Performance of Planar Floquet Codes with Majorana-Based Qubits](#), by Adam Paetzniak, Nicolas Delfosse, et al., PRX Quantum, January 2023 (15 pages).
- [1661] [Benchmarking the Planar Honeycomb Code](#), by Craig Gidney, Michael Newman, and Matt McEwen, arXiv, September 2022 (18 pages).
- [1662] [Planar Floquet Codes](#), by Christophe Vuillot, arXiv, December 2021 (16 pages).
- [1663] [A Pair Measurement Surface Code on Pentagons](#), by Craig Gidney, arXiv, October 2023 (15 pages).
- [1664] [Accommodating Fabrication Defects on Floquet Codes with Minimal Hardware Requirements](#), by Campbell McLauchlan, György P. Gehér, and Alexandra E. Moylett, arXiv, December 2024 (37 pages).
- [1665] [Quantum computation from dynamic automorphism codes](#), by Margarita Davydova, David Aasen, et al., arXiv, August 2024 (17 pages).
- [1666] [Dynamical quantum codes and logic gates on a lattice with sparse connectivity](#), by Dominic J. Williamson and Bence Hetényi, arXiv, October 2025 (31 pages).
- [1667] [Stairway Codes: Floquetifying Bivariate Bicycle Codes and Beyond](#), by Shoham Jacoby, Alex Retzker, and Fernando Pastawski, arXiv, February 2026 (18 pages).
- [1668] [Logical gates on Floquet codes via folds and twists](#), by Alexandra E. Moylett and Bhargavi Jonnadula, arXiv, January 2026 (49 pages).
- [1669] [Logical operations with a dynamical qubit in Floquet-Bacon-Shor code](#), by Xuandong Sun, Dapeng Yu, et al., arXiv, July 2025 (33 pages).
- [1670] [Quantum Reed-Muller Codes](#), by Andrew Steane, arXiv, August 1996 (9 pages).
- [1671] [Construction of the full logical Clifford group for high-rate quantum Reed-Muller codes using only transversal and fold-transversal gates](#), by Theerapat Tansuwanmont, Tim Chan, and Ryuji Takagi, arXiv, June 2026 (38 pages).
- [1672] [Continuous quantum error correction via quantum feedback control](#), by Charlene Ahn, Andrew C. Doherty, and Andrew J. Landahl, arXiv, October 2001 (12 pages).
- [1673] [A practical scheme for error control using feedback](#), by Mohan Sarovar, Charlene Ahn, Kurt Jacobs, and Gerard J. Milburn, arXiv, May 2004 (12 pages).
- [1674] [Measurement based estimator scheme for continuous quantum error correction](#), by Sangkha Borah, Jason Twamley, et al., arXiv, September 2022 (10 pages).
- [1675] [Emergent Non-Markovian Nonlinear Qubit From Collective Spin Interactions](#), by Gregory T. Carroll, Michael R. Geller, and Andre Erpenbeck, arXiv, August 2026 (9 pages).
- [1676] [Error Correction in Dynamical Codes](#), by Xiaozhen Fu and Daniel Gottesman, arXiv, March 2024 (42 pages).
- [1677] [Dynamical Logical Qubits in the Bacon-Shor Code](#), by M. Sohaib Alam and Eleanor Rieffel, arXiv, December 2024 (12 pages).
- [1678] [Autonomous quantum error correction beyond break-even and its metrological application](#), by Zhongchu Ni, Dapeng Yu, et al., arXiv, September 2025 (10 pages).
- [1679] [Encoding a qubit in an oscillator](#), by Daniel Gottesman, Alexei Kitaev, and John Preskill, arXiv, May 2001 (22 pages).
- [1680] [Performance analysis of GKP error correction](#), by Frederik K. Marqvorsen, Janus H. Wesenberg, Nikolaj T. Zinner, and Ulrik L. Andersen, arXiv, May 2025 (14 pages).
- [1681] [Hardware implementation of quantum stabilizers in superconducting circuits](#), by K. Dodge, B. L. T. Plourde, et al., arXiv, October 2023 (30 pages).
- [1682] [Quantum computing with rotation-symmetric bosonic codes](#), by Arne L. Grimsmo, Joshua Combes, and Ben Q. Baragiola, arXiv, September 2019 (31 pages).
- [1683] [Quantum spherical codes](#), by Shubham P. Jain, Victor V. Albert, et al., Nature Physics, May 2024  (preprint on arXiv).
- [1684] [Quantum control of a cat qubit with bit-flip times exceeding ten seconds](#), by Ulysse Réglade, Pierre Rouchon, Alain Sarlette, Mazyar Mirrahimi, Philippe Campagne-Ibarcq, Raphaël Lescanne, Sébastien Jezouin, and Zaki Leghtas, Nature, May 2024  (preprint on HAL).
- [1685] [Quantum error correction using squeezed Schrödinger cat states](#), by David S. Schlegel, Fabrizio Minganti, and Vincenzo Savona, arXiv, September 2022 (19 pages).
- [1686] [Protecting quantum entanglement between error-corrected logical qubits](#), by Weizhou Cai, Luyan Sun, et al., arXiv, February 2023 (26 pages).
- [1687] [The 2T-qutrit, a two-mode bosonic qutrit](#), by Aurélie Denys and Anthony Leverrier, arXiv, May 2023 (24 pages).
- [1688] [Quantum error-correcting codes with a covariant encoding](#), by Aurélie Denys and Anthony Leverrier, arXiv, February 2025 (13 pages).
- [1689] [Autonomous error correction of a single logical qubit using two transmons](#), by Ziqian Li, David I. Schuster, et al., arXiv, February 2023 (20 pages).
- [1690] [Tailoring quantum error correction to spin qubits](#), by Bence Hetényi and James R. Wootton, arXiv, March 2024 (25 pages).

- [1691] [Quantum error correction with the Gottesman-Kitaev-Preskill code](#), by Arne Grimsmo et al., PRX Quantum, June 2021 (20 pages).
- [1692] [Increasing error tolerance in quantum computers with dynamic bias arrangement](#), by Hector Bombín, Naomi Nickerson, Jordan Sullivan, et al., arXiv, March 2023 (11 pages).
- [1693] [Modular decoding: parallelizable real-time decoding for quantum computers](#), by Héctor Bombín, Chris Dawson, Ye-Hua Liu, Naomi Nickerson, Fernando Pastawski, and Sam Roberts, arXiv, March 2023 (2 pages).
- [1694] [Why and what is the future of the topological qubit?](#), by Chetan Nayak, 2022.
- [1695] [High threshold codes for neutral atom qubits with biased erasure errors](#), by Kaavya Sahay, Shruti Puri, et al., arXiv, October 2023 (8 pages).
- [1696] [Quantum Error Correction with Metastable States of Trapped Ions Using Erasure Conversion](#), by Mingyu Kang, Wesley C. Campbell, and Kenneth R. Brown, arXiv, June 2023 (20 pages).
- [1697] [Error mitigation, optimization, and extrapolation on a trapped ion testbed](#), by Oliver G. Maupin, Peter J. Love, et al., arXiv, September 2024 (21 pages).
- [1698] [Erasure conversion in Majorana qubits via local quasiparticle detection](#), by Abhijeet Alase, Kevin D. Stubbs, Barry C. Sanders, and David L. Feder, arXiv, December 2024 (23 pages).
- [1699] [Cleaner magic states with hook injection](#), by Craig Gidney, arXiv, February 2023 (12 pages).
- [1700] [Experiments with the 4D Surface Code on a QCCD Quantum Computer](#), by Noah Berthussen, David Hayes, et al., arXiv, December 2024 (12 pages).
- [1701] [Fault-tolerant error correction with the gauge color code](#), by Benjamin J Brown, Daniel Loss, and Yiannakis K Pachos, Nature Communications, July 2016.
- [1702] [Asymptotically good quantum and locally testable classical LDPC codes](#), by Pavel Panteleev and Gleb Kalachev, ACM Proceedings STOC 2022, June 2022.
- [1703] [Layer codes](#), by Dominic J. Williamson and Nouédyne Baspin, Nature Communications, November 2024.
- [1704] [Single-shot quantum error correction with the three-dimensional subsystem toric code](#), by Aleksander Kubica and Michael Vasmer, Nature Communications, October 2022.
- [1705] [High-performance fault-tolerant quantum computing with many-hypercube codes](#), by Hayato Goto, Science, September 2024.
- [1706] [Single-shot error correction with the gauge color code](#), by Hector Bombin, Physical Review X, September 2015.
- [1707] [Microsoft advances quantum error correction with a family of novel four-dimensional codes - Microsoft Azure Quantum Blog](#), by Krysta Svore, Microsoft Azure Quantum Blog, June 2025.
- [1708] [A Topologically Fault-Tolerant Quantum Computer with Four Dimensional Geometric Codes](#), by David Aasen, Matthew B. Hastings, Krysta M. Svore, et al., arXiv, June 2025 (29 pages).
- [1709] [Programmable quantum simulations of spin systems with trapped ions](#), by Christopher Monroe et al., Reviews of Modern Physics, April 2021 (preprint on arXiv).
- [1710] [ERASER: Towards Adaptive Leakage Suppression for Fault-Tolerant Quantum Computing](#), by Suhas Vittal, Poulami Das, and Moinuddin Qureshi, arXiv, September 2023 (17 pages).
- [1711] [Quantum Error Correction resilient against Atom Loss](#), by Hugo Perrin, Sven Jandura, and Guido Pupillo, arXiv, March 2025 (26 pages).
- [1712] [Fast erasure decoder for hypergraph product codes](#), by Nicholas Connolly, Vivien Londe, Anthony Leverrier, and Nicolas Delfosse, arXiv, August 2024 (11 pages).
- [1713] [AI-Enabled Decoding of Qubit Loss for Quantum Error-Correcting Codes](#), by Yuqing Wang, Xiaotian Nie, Jiale Dai, Zhongyi Ni, Tao Zhang, Hui Zhai, and Linghui Chen, arXiv, May 2026 (12 pages).
- [1714] [Circuit-based leakage-to-erasure conversion in a neutral atom quantum processor](#), by Matthew N. H. Chow, Ivan H. Deutsch, Yuan-Yu Jau, et al., arXiv, December 2024 (14 pages).
- [1715] [High-fidelity gates with mid-circuit erasure conversion in a metastable neutral atom qubit](#), by Shuo Ma, Guido Pupillo, Shruti Puri, Jeff D. Thompson, et al., arXiv, May 2023 (17 pages).
- [1716] [Erasure conversion in a high-fidelity Rydberg quantum simulator](#), by Pascal Scholl, Manuel Endres, et al., arXiv, October 2023 (21 pages).
- [1717] [Loss-biased fault-tolerant quantum error correction](#), by Laura Pecorari, Gavin K. Brennen, Stanimir S. Kondov, and Guido Pupillo, arXiv, April 2026 (9 pages).
- [1718] [High-fidelity gates and mid-circuit erasure conversion in an atomic qubit](#), by Shuo Ma, Guido Pupillo, Shruti Puri, Jeff D. Thompson, et al., Nature, October 2023 (preprint on arXiv).
- [1719] [Quantum LDPC codes for erasure-biased atomic quantum processors](#), by Laura Pecorari and Guido Pupillo, arXiv, February 2025 (12 pages).
- [1720] [Demonstrating a superconducting dual-rail cavity qubit with erasure-detected logical measurements](#), by Kevin S. Chou, Robert J. Schoelkopf, et al., arXiv, October 2023 (38 pages).
- [1721] [Demonstrating a long-coherence dual-rail erasure qubit using tunable transmons](#), by Harry Levine, Arbel Haim, et al., arXiv, March 2024 (13 pages).
- [1722] [Optimizing quantum error correction protocols with erasure qubits](#), by Shouzen Gu, Yotam Vaknin, Alex Retzker, and Aleksander Kubica, arXiv, August 2024 (6 pages).
- [1723] [Surface Code with Imperfect Erasure Checks](#), by Kathleen Chang, Shruti Puri, et al., arXiv, August 2024 (19 pages).
- [1724] [Flag fault-tolerant error correction with arbitrary distance codes](#), by Christopher Chamberland and Michael E. Beverland, arXiv, January 2018 (29 pages).
- [1725] [Optimal control with flag qubits](#), by Liang-Xu Xie, Lui Zuccherelli de Paula, Weizhou Cai, Qing-Xuan Jie, Luyan Sun, Chang-Ling Zou, Guang-Can Guo, Zi-Jie Chen, and Xu-Bo Zou, arXiv, March 2026 (8 pages).
- [1726] [Medusa: Detecting and Removing Failures for Scalable Quantum Computing](#), by Karoliina Oksanen, Quan Hoang, and Alexandru Paler, arXiv, July 2026 (7 pages).
- [1727] [Extending the lifetime of a quantum bit with error correction in superconducting circuits](#), by Nissim Ofek et al., 2016 (5 pages).
- [1728] [Real-time quantum error correction beyond break-even](#), by V. V. Sivak, Robert J. Schoelkopf, Michel H. Devoret, et al., Nature, March 2023 (preprint on arXiv).
- [1729] [Accelerating Fault-Tolerant Quantum Computation with Good qLDPC Codes](#), by Guo Zhang, Yuanye Zhu, and Ying Li, arXiv, October 2025 (57 pages).
- [1730] [Dynamic Concatenation of Quantum Error Correction in Integrated Quantum Computing Architecture](#), by Ilkwon Sohn, 2019 (7 pages).
- [1731] [Learning to Concatenate Quantum Codes](#), by Nico Meyer, Christopher Mutschler, Dominik Seuß, Andreas Maier, and Daniel D. Scherer, arXiv, April 2026 (7 pages).
- [1732] [Tutorial : general introduction to Fault-Tolerant Quantum Computing](#), by Keisuke Fujii, July 2023 (49 slides).
- [1733] [Fault-Tolerant Quantum Computation With Constant Error Rate](#), by Dorit Aharonov and Michael Ben-Or, arXiv, June 1999 (63 pages).
- [1734] [Optimizing resource efficiencies for scalable full-stack quantum computers](#), by Marco Fellous-Asiani, Jing Hao Chai,

- Yvain Thonnart, Hui Khoon Ng, Robert S. Whitney, and Alexia Auffèves, arXiv, October 2023 (41 pages).
- [1735] [Fault-tolerant quantum error correction on near-term quantum processors using flag and bridge qubits](#), by Lao Lingling, 2020 (12 pages).
- [1736] [Fault-tolerant quantum error correction for Steane's seven-qubit color code with few or no extra qubits](#), by Ben W. Reichardt, arXiv, April 2018 (11 pages).
- [1737] [Overhead analysis of universal concatenated quantum codes](#), by Christopher Chamberland, Tomas Jochym-O'Connor, and Raymond Laflamme, arXiv, June 2017 (25 pages).
- [1738] [Tailoring Fault-Tolerance to Quantum Algorithms](#), by Zhuangzhuang Chen and Narayanan Rengaswamy, arXiv, April 2024 (19 pages).
- [1739] [Fault-tolerant operation of a logical qubit in a diamond quantum processor](#), by M. H. Abobeih, Barbara M. Terhal, T. H. Taminiau, et al., Nature, May 2022.
- [1740] [Architectures for Heterogeneous Quantum Error Correction Codes](#), by Samuel Stein, Ang Li, et al., arXiv, November 2024 (13 pages).
- [1741] [Compilation Pipeline for Predicting Algorithmic Break-Even in an Early-Fault-Tolerant Surface Code Architecture](#), by Tianyi Hao, Joseph Sullivan, Sivaprasad Omanakuttan, Michael A. Perlin, and Ruslan Shaydulin, arXiv, November 2025 (3 pages).
- [1742] [Quasilinear Equivalence Checking for Detector Error Models](#), by Mathys Rennela, arXiv, June 2026 (19 pages).
- [1743] [LightStim: A Framework for QEC Protocol Evaluation and Prototyping with Automated DEM Construction](#), by Xiang Fang, Ming Wang, Yue Wu, Sharanya Prabhu, Dean Tullsen, Narasinga Rao Miniskar, Frank Mueller, Travis Humble, and Yufei Ding, arXiv, June 2026 (18 pages).
- [1744] [Sharpening Worst-Case Error Assessment for Fault-Tolerant Quantum Computing: Fidelity and Its Deviation](#), by Kyoungcho Cho, Ilkwon Sohn, Yongsoo Hwang, and Jeongho Bang, arXiv, March 2026 (48 pages).
- [1745] [Performance Limits of Fault-Tolerant Quantum Error Correction Schemes](#), by Lorenzo Valentini, Diego Forlivesi, and Marco Chiani, arXiv, May 2026 (15 pages).
- [1746] [Boundaries of Acceptable Defectiveness: Redefining Surface Code Robustness under Heterogeneous Noise](#), by Jacob S. Palmer and Kaitlin N. Smith, arXiv, March 2026 (12 pages).
- [1747] [Surface code error correction on a defective lattice](#), by Shota Nagayama, Austin G Fowler, Dominic Horsman, Simon J Devitt, and Rodney Van Meter, New Journal of Physics, February 2017 (29 pages).
- [1748] [Remote Entanglement in Lattice Surgery: To Distill, or Not to Distill](#), by Sitong Liu, John Stack, Ke Sun, Roel Van Beeumen, Inder Monga, Katherine Klymko, Kenneth R. Brown, and Erhan Saglamyurek, arXiv, May 2026 (19 pages).
- [1749] [Towards Predictive Quantum Algorithmic Performance: Modeling Time-Correlated Noise at Scale](#), by Amit Jamadagni, Gregory Quiroz, and Eugene Dumitrescu, arXiv, March 2026 (13 pages).
- [1750] [Decoherence benchmarking of superconducting qubits](#), by Jonathan J. Burnett, Jonas Bylander, et al., npj Quantum Information, June 2019.
- [1751] [X-Z Round Scheduling for the Surface Code with Defects under Biased Noise](#), by Lakshika Rathi, Pau Escofet, Joshua Vizlai, and Margaret Martonosi, arXiv, August 2026 (11 pages).
- [1752] [Simulating Quantum Error Correction beyond Pauli Stochastic Errors](#), by Jordan Hines, Corey Ostrove, Kenneth Rudinger, Stefan Seritan, Kevin Young, Robin Blume-Kohout, and Timothy Proctor, arXiv, March 2026 (11 pages).
- [1753] [Comparing Experiments to the Fault-Tolerance Threshold](#), by Richard Kueng, David M. Long, Andrew C. Doherty, and Steven T. Flammia, arXiv, July 2016 (18 pages).
- [1754] [Dynamical Decoupling of Open Quantum Systems](#), by Lorenza Viola, Emanuel Knill, and Seth Lloyd, Physical Review Letters, March 1999 (preprint on arXiv).
- [1755] [Fault-Tolerant Quantum Error Detection](#), by Norbert M. Linke, Mauricio Gutierrez, Kevin A. Landsman, Caroline Figgatt, Shantanu Debnath, Kenneth R. Brown, and Christopher Monroe, Science Advances, October 2017 (7).
- [1756] [Real-Time Quantum Error Correction System Stack: Architecture, Algorithms, and Engineering Practice](#), by Yaojian Chen, Dingshun Lv, et al., arXiv, 2026 (55 pages).
- [1757] [Adaptive Estimation of Drifting Noise in Quantum Error Correction](#), by Devansh Bhardwaj, Evangelia Takou, Yingjia Lin, and Kenneth R. Brown, arXiv, November 2025 (20 pages).
- [1758] [How to choose a decoder for a fault-tolerant quantum computer? The speed vs accuracy trade-off](#), by Nicolas Delfosse, Andres Paz, Alexander Vaschillo, and Krysta M. Svore, arXiv, October 2023 (19 pages).
- [1759] [Hardness results for decoding the surface code with Pauli noise](#), by Alex Fischer and Akimasa Miyake, arXiv, November 2024 (44 pages).
- [1760] [Splitting decoders for correcting hypergraph faults](#), by Nicolas Delfosse, Adam Paetznick, Jeongwan Haah, and Matthew B. Hastings, arXiv, September 2023 (12 pages).
- [1761] [The color code, the surface code, and the transversal CNOT: NP-hardness of minimum-weight decoding](#), by Shouzen Gu, Lily Wang, and Aleksander Kubica, arXiv, March 2026 (18 pages).
- [1762] [Doubly-polylog-time-overhead fault-tolerant quantum computation by a polylog-time parallel minimum-weight perfect matching decoder](#), by Yugo Takada and Hayata Yamasaki, arXiv, March 2025 (35 pages).
- [1763] [Neural Minimum Weight Perfect Matching for Quantum Error Codes](#), by Yotam Peled, David Zenati, and Eliya Nachmani, arXiv, June 2026 (16 pages).
- [1764] [A matching decoder for bivariate bicycle codes](#), by Kaavya Sahay, Dominic J. Williamson, and Benjamin J. Brown, arXiv, February 2026 (23 pages).
- [1765] [Minimum-Weight Parity Factor Decoder for Quantum Error Correction](#), by Yue Wu, Binghong Li, Kathleen Chang, Shruti Puri, and Lin Zhong, arXiv, August 2025 (44 pages).
- [1766] [Tesseract: A Search-Based Decoder for Quantum Error Correction](#), by Laleh Aghababae Beni, Oscar Higgott, and Noah Shutty, arXiv, March 2025 (13 pages).
- [1767] [Accelerating the Tesseract Decoder for Quantum Error Correction](#), by Dragana Grbic, Laleh Aghababae Beni, and Noah Shutty, arXiv, February 2026 (22 pages).
- [1768] [An interpretation of Union-Find Decoder on Weighted Graphs](#), by Yue Wu, Namitha Liyanage, and Lin Zhong, arXiv, November 2022 (20 pages).
- [1769] [Breadth-first graph traversal union-find decoder](#), by Matthias C. Löbl, Susan X. Chen, Stefano Paesani, and Anders S. Sørensen, arXiv, July 2024 (17 pages).
- [1770] [Proof of a finite threshold for the union-find decoder](#), by Satoshi Yoshida, Ethan Lake, and Hayata Yamasaki, arXiv, May 2026 (20 pages).
- [1771] [Automorphism Ensemble Decoding of Quantum LDPC Codes](#), by Stergios Koutsoumpas, Hasan Sayginel, Mark Webster, and Dan E Browne, arXiv, March 2025 (10 pages).
- [1772] [Degeneracy Cutting: A Local and Efficient Post-Processing for Belief Propagation Decoding of Quantum Low-Density Parity-Check Codes](#), by Kento Tsubouchi, Hayata Yamasaki, and Shiro Tamiya, arXiv, October 2025 (17 pages).
- [1773] [Beam search decoder for quantum LDPC codes](#), by Min Ye, Dave Wecker, and Nicolas Delfosse, arXiv, December 2025 (11 pages).

- [1774] [Achieving Thresholds via Standalone Belief Propagation on Surface Codes](#), by Pedro Hack, Luca Menti, Francisco Lazaro, and Alexandru Paler, arXiv, March 2026 (9 pages).
- [1775] [Learning-Based List Sequential Belief Propagation Decoding of Quantum LDPC Codes](#), by Mohsen Moradi, Taejoon Kim, and Remi A. Chou, arXiv, June 2026 (13 pages).
- [1776] [Degenerate Quantum LDPC Codes With Good Finite Length Performance](#), by Pavel Panteleev and Gleb Kalachev, arXiv, November 2021 (21 pages).
- [1777] [Decoding Across the Quantum LDPC Code Landscape](#), by Joschka Roffe, David R. White, Simon Burton, and Earl T. Campbell, arXiv, December 2020 (15 pages).
- [1778] [Affine Subcode Ensemble Decoding for Degeneracy-Aware Quantum Error Correction](#), by Leo Wursthorn, Jonathan Mandelbaum, Sisi Miao, Hedongliang Liu, Holger Jäkel, Stergios Koutsoumpas, and Laurent Schmalen, arXiv, May 2026 (6 pages).
- [1779] [Fully convolutional 3D neural network decoders for surface codes with syndrome circuit noise](#), by Spiro Gicev, Lloyd C. L. Hollenberg, and Muhammad Usman, arXiv, June 2025 (34 pages).
- [1780] [Generative Decoding for Quantum Error-correcting Codes](#), by Hanyan Cao, Feng Pan, Dongyang Feng, Yijia Wang, and Pan Zhang, arXiv, March 2025 (12 pages).
- [1781] [Deep Quantum Error Correction](#), by Yoni Choukroun and Lior Wolf, arXiv, December 2023 (11 pages).
- [1782] [Machine Learning Approaches to Decoding Topological Quantum Codes](#), by Changwon Lee, Tak Hur, Jeongwoo Jae, and Daniel K. Park, arXiv, August 2026 (42 pages).
- [1783] [Reinforcement Learning for Enhanced Advanced QEC Architecture Decoding](#), by Yidong Zhou, Lingyi Kong, Yifeng Peng, and Zhiding Liang, arXiv, January 2026 (7 pages).
- [1784] [Efficient foundation decoders for fault-tolerant quantum computing](#), by Ge Yan, Shanchuan Li, Shiyi Xiao, Pengyue Ma, Hanyan Cao, Feng Pan, and Yuxuan Du, arXiv, June 2026 (32 pages).
- [1785] [Learning to Decode in Parallel: Self-Coordinating Neural Network for Real-Time Quantum Error Correction](#), by Kai Zhang, Zhengzhong Yi, Shaojun Guo, Linghang Kong, Situ Wang, Xiaoyu Zhan, Tan He, Weiping Lin, Tao Jiang, Dongxin Gao, Yiming Zhang, Fangming Liu, Fang Zhang, Zhengfeng Ji, Fusheng Chen, and Jianxin Chen, arXiv, January 2026 (25 pages).
- [1786] [A Mixture of Experts Vision Transformer for High-Fidelity Surface Code Decoding](#), by Hoang Viet Nguyen, Manh Hung Nguyen, Hoang Ta, Van Khu Vu, and Yeow Meng Chee, arXiv, January 2026 (16 pages).
- [1787] [HyperNQ: A Hypergraph Neural Network Decoder for Quantum LDPC Codes](#), by Ameysa S. Bhave, Navnil Choudhury, and Kanad Basu, arXiv, November 2025 (6 pages).
- [1788] [AlphaSyndrome: Tackling the Syndrome Measurement Circuit Scheduling Problem for QEC Codes](#), by Yuhao Liu, Shuohao Ping, Junyu Zhou, Ethan Decker, Justin Kalloor, Mathias Weiden, Kean Chen, Yunong Shi, Ali Javadi-Abhari, Costin Iancu, and Gushu Li, arXiv, February 2026 (17 pages).
- [1789] [Optimal Decoding with the Worm](#), by Zac Tobias, Nikolas P. Breuckmann, and Benedikt Placke, arXiv, March 2026 (34 pages).
- [1790] [A scalable and real-time neural decoder for topological quantum codes](#), by Andrew W. Senior, Hartmut Neven, Pushmeet Kohli, Johannes Bausch, et al., arXiv, December 2025 (48 pages).
- [1791] [Rethink the Role of Neural Decoders in Quantum Error Correction](#), by Ge Yan, Shanchuan Li, and Yuxuan Du, arXiv, May 2026 (33 pages).
- [1792] [Neural Decoders for Universal Quantum Algorithms](#), by J. Pablo Bonilla Ataides, Andi Gu, Susanne F. Yelin, and Mikhail D. Lukin, arXiv, September 2025 (12 pages).
- [1793] [Scalable Neural Decoders for Practical Fault-Tolerant Quantum Computation](#), by Andi Gu, J. Pablo Bonilla Ataides, Mikhail D. Lukin, and Susanne F. Yelin, arXiv, April 2026 (18 pages).
- [1794] [SymBreak: Mitigating Quantum Degeneracy Issues in QLDPC Code Decoders by Breaking Symmetry](#), by Keyi Yin, Travis Humble, Yufei Ding, et al., arXiv, December 2024 (14 pages).
- [1795] [SOME: Symmetric One-Hot Matching Elector – A Lightweight Microsecond Decoder for Quantum Error Correction](#), by Xinyi Guo, Takashi Sato, et al., arXiv, July 2025 (9 pages).
- [1796] [Adaptive Window Decoding based on Spatiotemporal Complementary Gap](#), by Moeto Mishima, Riki Toshio, Kaito Kishi, Jun Fujisaki, Hirotaka Oshima, Shintaro Sato, and Keisuke Fujii, arXiv, May 2026 (22 pages).
- [1797] [Quantum Maxwell Erasure Decoder for qLDPC codes](#), by Bruno Costa Alves Freire, François-Marie Le Régent, and Anthony Leverrier, arXiv, January 2026 (8 pages).
- [1798] [LEAD: A Local Ensemble-Assisted Parallel Decoding Framework for Quantum Tanner Codes](#), by Zhuo-Yan Xiao, Sha Shi, Chen-Peng Huang, Dong-Sheng Wang, and Yun-Jiang Wang, arXiv, May 2026 (10 pages).
- [1799] [Hierarchical quantum decoders](#), by Nirupam Basak, Ankith Mohan, Andrew Tanggara, Tobias Haug, Goutam Paul, and Kishor Bharti, arXiv, January 2026 (17 pages).
- [1800] [Near-optimal decoding algorithm for color codes using Population Annealing](#), by Fernando Martínez-García, Pedro Parrado-Rodríguez, et al., arXiv, May 2024 (11 pages).
- [1801] [Decoding Quantum LDPC Codes Using Graph Neural Networks](#), by Vukan Ninkovic, Alexandre Graell i Amat, et al., arXiv, August 2024 (6 pages).
- [1802] [Efficient and Universal Neural-Network Decoder for Stabilizer-Based Quantum Error Correction](#), by Gengyuan Hu, Han-Sen Zhong, et al., arXiv, June 2025 (36 pages).
- [1803] [A quantum annealing approach to the minimum distance problem of quantum codes](#), by Refat Ismail, Ashish Kakkar, and Anatoly Dymarsky, arXiv, July 2024 (20 pages).
- [1804] [Correcting a noisy quantum computer using a quantum computer](#), by Pan Zhang, arXiv, June 2025 (13 pages).
- [1805] [Quantum Approximate Optimization for Decoding of Low-Density Parity-Check Codes](#), by Krishnakanta Barik and Goutam Paul, arXiv, February 2026 (8 pages).
- [1806] [Adaptive Aborting Schemes for Quantum Error Correction Decoding](#), by Sanidhay Bhambay, Prakash Murali, Neil Walton, and Thirupathaiah Vasantam, arXiv, February 2026 (31 pages).
- [1807] [Towards Ultra-High-Rate Quantum Error Correction with Reconfigurable Atom Arrays](#), by Chen Zhao, Casey Duckering, Andi Gu, Nishad Maskara, and Hengyun Zhou, arXiv, April 2026 (21 pages).
- [1808] [Multistage Rewinding Decoder for QLDPC Codes](#), by Milad Taghipour, Dimitris Chytas, and Bane Vasić, arXiv, August 2026 (9 pages).
- [1809] [Simplified circuit-level decoding using Knill error correction](#), by Ewan Murphy, Subhayan Sahu, and Michael Vasmer, arXiv, April 2026 (7 pages).
- [1810] [Platform Architecture for Tight Coupling of High-Performance Computing with Quantum Processors](#), by Shane A. Caldwell, Travis Humble, et al., arXiv, October 2025 (43 pages).
- [1811] [Real-time decoding of quantum error correction codes using high-performance computing](#), by Lingling Lao, Qiang Wang, Yuanqi Liu, Yantong Liu, Haowen Wang, Yitao Chen, Yankang Zhao, Zhenwei Wu, Wei Zhang, Yong Dong, Yingwen Liu, Mingche Lai, and Junjie Wu, arXiv, August 2026 (11 pages).
- [1812] [Soft decoding for quantum LDPC codes with experimental validation](#), by Arda Aydin, Edwin Tham, Nicolas Delfosse, and Min Ye, arXiv, September 2026 (9 pages).

- [1813] GPU-Accelerated Syndrome Decoding for Quantum LDPC Codes below the 63 μ s Latency Threshold, by Oscar Ferraz, Vitor Silva, et al., arXiv, August 2025 (7 pages).
- [1814] Sparse Blossom: correcting a million errors per core second with minimum-weight matching, by Oscar Higgott and Craig Gidney, arXiv, January 2025 (37 pages).
- [1815] Decoding quantum color codes with MaxSAT, by Lucas Berent, Lukas Burgholzer, Peter-Jan H.S. Derks, Jens Eisert, and Robert Wille, arXiv, October 2024 (17 pages).
- [1816] Enabling Full-Stack Quantum Computing with Changeable Error-Corrected Qubits, by Anbang Wu, Keyi Yin, Andrew W. Cross, Ang Li, and Yufei Ding, arXiv, October 2023 (13 pages).
- [1817] Minimum Weight Decoding in the Colour Code is NP-hard, by Mark Walters and Mark L. Turner, arXiv, March 2026 (18 pages).
- [1818] FPGA-tailored algorithms for real-time decoding of quantum LDPC codes, by Satvik Maurya, Thilo Maurer, Markus Böhler, Drew Vandeth, and Michael E. Beverland, arXiv, January 2026 (63 pages).
- [1819] Real-time Surface-Code Error Correction Using an FPGA-based Neural-Network Decoder, by Xiaohan Yang, Xuan-dong Sun, Zhiyi Wu, Jiawei Zhang, Ji Jiang, Xiayu Linpeng, Yuxuan Zhou, Ji Chu, Jingjing Niu, Youpeng Zhong, Song Liu, and Dapeng Yu, arXiv, May 2026 (27 pages).
- [1820] A Scalable FPGA Architecture for Real-Time Decoding of Quantum LDPC Codes Using GARI, by Daniel Bascónes, Arshpreet Singh Maan, Valentin Savin, and Francisco Garcia-Herrero, arXiv, May 2026 (11 pages).
- [1821] Zero-G: A Pre-Decoder-Aware Decoder for Quantum Error Correction, by Peter Wegmann, Theofilos Augoustis, Aleksandra Świerkowska, Emmanouil Giortamis, and Pramod Bhatotia, arXiv, August 2026.
- [1822] QAdapt: A Noise-Adaptive Neural Pre-Decoding Framework for Quantum Error Correction, by Ran Miao, Rui Luo, Xiaohan Shan, and Xiaoming Sun, arXiv, July 2026 (11 pages).
- [1823] Pinball: A Cryogenic Predecoder for Surface Code Decoding Under Circuit-Level Noise, by Alexander Knapen, Guanchen Tao, Jacob Mack, Tomas Bruno, Mehdi Saligane, Dennis Sylvester, Qirui Zhang, and Gokul Subramanian Ravi, arXiv, December 2025 (17 pages).
- [1824] A local pre-decoder to reduce the bandwidth and latency of quantum error correction, by Samuel C. Smith, Benjamin J. Brown, and Stephen D. Bartlett, arXiv, September 2022 (12 pages).
- [1825] What determines the ultimate precision of a quantum computer?, by Xavier Waintal, arXiv, April 2019 (6 pages).
- [1826] Stabilizer Formalism for Operator Quantum Error Correction, by David Poulin, PRL, 2005 (5 pages) (preprint on arXiv).
- [1827] Operator Quantum Error Correcting Subsystems for Self-Correcting Quantum Memories, by Dave Bacon, arXiv, 2006 (17 pages).
- [1828] Fault-Tolerant Operation of a Quantum Error-Correction Code, by Laird Egan, Christopher Monroe, et al., arXiv, January 2021 (23 pages).
- [1829] Fault-Tolerant Quantum Computing with Trapped Ions: The Walking Cat Architecture, by Felix Tripier, Woo Chang Chung, Jacob Young, Safwan Alam, Bryce Bjork, Aharon Brodutch, Finn Lasse Buessen, Nolan J. Coble, Thomas Dellaert, Dmitri Maslov, Martin Roetteler, Edwin Tham, Mark Webster, Min Ye, John Gamble, Andrii Maksymov, J. P. Marceaux, and Nicolas Delfosse, arXiv, April 2026 (110 pages).
- [1830] Suppressing quantum errors by scaling a surface code logical qubit, by Google Quantum AI, Nature, February 2023.
- [1831] Quantum error correction below the surface code threshold, by Google Quantum AI and Collaborators, Nature, December 2024.
- [1832] Constant Overhead Quantum Fault Tolerance with Quantum Expander Codes, by Omar Fawzi et al., Communications of the ACM, 2020 (9 pages) (preprint on HAL).
- [1833] Time-Efficient Constant-Space-Overhead Fault-Tolerant Quantum Computation, by Hayata Yamasaki and Masato Koashi, Nature Physics, January 2024.
- [1834] Polylog-time- and constant-space-overhead fault-tolerant quantum computation with quantum low-density parity-check codes, by Shiro Tamiya, Masato Koashi, and Hayata Yamasaki, arXiv, December 2024 (63 pages).
- [1835] Quantum fault tolerance with constant-space and logarithmic-time overheads, by Quynh T. Nguyen and Christopher A. Pattison, arXiv, November 2024 (119 pages).
- [1836] Purely-logarithmic-time- and constant-space-overhead fault-tolerant quantum computation, by Zhengyi Han and Zi-Wen Liu, arXiv, September 2026 (71 pages).
- [1837] Improved Bounds for Rounding Errors in Quantum Circuit Simulators, by Jonas Klamroth et al., Software Engineering 2025, February 2025 (15 pages).
- [1838] Fault-Tolerant Constant-Depth Clifford Gates on Toric Codes, by Alexandre Guernut and Christophe Vuillot, arXiv, November 2024 (35 pages).
- [1839] Fault-tolerant quantum computation cannot be achieved with constant spacetime overhead, by Kishor Bharti, Tobias Haug, and Andrew Tanggara, arXiv, August 2026 (40 pages).
- [1840] No-Go Theorems for Universal Entanglement Purification, by Allen Zang et al., Physical Review Letters, May 2025 (1).
- [1841] Taming the Instruction Bandwidth of Quantum Computers via Hardware Managed Error Correction, by Swamit Tannu et al., 2017 (13 pages).
- [1842] Parallel window decoding enables scalable fault tolerant quantum computation, by Luka Skoric, Earl T. Campbell, et al., arXiv, February 2023 (12 pages).
- [1843] Real-Time Decoding for Fault-Tolerant Quantum Computing: Progress, Challenges and Outlook, by Francesco Battistel et al., Nano Futures, August 2023 (15 pages) (preprint on arXiv).
- [1844] No More Hooks in the Surface Code: Distance-Preserving Syndrome Extraction for Arbitrary Layouts at Minimum Depth, by Yuga Hirai, Shota Ikari, Yosuke Ueno, and Yasunari Suzuki, arXiv, June 2026 (18 pages).
- [1845] Surface code off-the-hook: diagonal syndrome-extraction scheduling, by Gilad Kishony and Austin Fowler, arXiv, February 2026 (11 pages).
- [1846] Color code off-the-hook: avoiding hook errors with a single auxiliary per plaquette, by Gilad Kishony and Austin Fowler, arXiv, March 2026 (6 pages).
- [1847] Cyclic Hypergraph Product Code, by Arda Aydin, Nicolas Delfosse, and Edwin Tham, arXiv, November 2025 (10 pages).
- [1848] Full Extractors for Logical Processing in Hypergraph Product Codes, by John Blue, Zhiyang He, Hengyun Zhou, and Isaac L. Chuang, arXiv, June 2026 (30 pages).
- [1849] Realizing Repeated Quantum Error Correction in a Distance-Three Surface Code, by Sebastian Krinner, Alexandre Blais, Christopher Eichler, Andreas Wallraff, et al., Nature, May 2022 (preprint on arXiv).
- [1850] Demonstration of fault-tolerant universal quantum gate operations, by Lukas Postler, Thomas Monz, et al., Nature, May 2022 (14 pages) (preprint on arXiv).
- [1851] Realization of an Error-Correcting Surface Code with Superconducting Qubits, by Youwei Zhao, Jian-Wei Pan, et al., arXiv, January 2022 (10 pages).
- [1852] Realization of Real-Time Fault-Tolerant Quantum Error Correction, by C. Ryan-Anderson, PRX, December 2021 (29 pages).

- [1853] [Implementing Fault-tolerant Entangling Gates on the Five-qubit Code and the Color Code](#), by C. Ryan-Anderson, David Hayes, et al., arXiv, August 2022 (17 pages).
- [1854] [Demonstration of Fault-Tolerant Steane Quantum Error Correction](#), by Lukas Postler et al., PRX Quantum, December 2023 (19 pages).
- [1855] [Demonstrating Bayesian Quantum Phase Estimation with Quantum Error Detection](#), by Kentaro Yamamoto, David Muñoz Ramo, et al., arXiv, September 2023 (16 pages).
- [1856] [Logical quantum processor based on reconfigurable atom arrays](#), by Dolev Bluvstein, Mikhail D. Lukin, et al., Nature, December 2023 (28 pages).
- [1857] [Entangling four logical qubits beyond break-even in a non-local code](#), by Yifan Hong, Elijah Durso-Sabina, David Hayes, and Andrew Lucas, arXiv, August 2024 (13 pages).
- [1858] [Demonstration of quantum computation and error correction with a tesseract code](#), by Ben W. Reichardt, Matt Zanner, et al., arXiv, December 2024 (12 pages).
- [1859] [Demonstration of logical qubits and repeated error correction with better-than-physical error rates](#), by A. Paetznick, M. P. da Silva, et al., arXiv, November 2024 (13 pages).
- [1860] [Demonstration of a Logical Architecture Uniting Motion and In-Plane Entanglement: Shor's Algorithm, Constant-Depth CNOT Ladder, and Many-Hypercube Code](#), by Rich Rines, Ilya Vinogradov, et al., arXiv, September 2025 (17 pages).
- [1861] [Fault-tolerant execution of error-corrected quantum algorithms](#), by Michael A. Perlin, Ruslan Shaydulin, et al., arXiv, March 2026 (32 pages).
- [1862] [Demonstration of 30 Logical Qubits on SqaLe](#), by Pranav Gokhale, Inflection blog, September 2026.
- [1863] [Fault-Tolerant One-Bit Addition with the Smallest Interesting Colour Code](#), by Yang Wang Ben Criger et al., arXiv, September 2023 (11 pages).
- [1864] [Realizing Lattice Surgery on Two Distance-Three Repetition Codes with Superconducting Qubits](#), by Ilya Besedin, Markus Müller, Andreas Wallraff, et al., arXiv, January 2025 (19 pages).
- [1865] [Exponential suppression of bit or phase errors with cyclic error correction](#), by Zijun Chen et al., Nature, February 2021 (6 pages).
- [1866] [Suppressing quantum errors by scaling a surface code logical qubit](#), by Rajeev Acharya, Igor Aleiner, et al., arXiv, July 2022 (6 pages).
- [1867] [QuTech and Fujitsu realise the fault-tolerant operation of a qubit](#), by QuTech, May 2022.
- [1868] [Fault-Tolerant Computing with Single Qudit Encoding](#), by Matteo Mezzadri, Alessandro Chiesa, Luca Lepori, and Stefano Carretta, arXiv, December 2023 (8 pages).
- [1869] [A Converse for Fault-tolerant Quantum Computation](#), by Uthirakalyani G, Anuj K. Nayak, and Avhishek Chatterjee, arXiv, August 2023 (14 pages).
- [1870] [Qudit vs. Qubit: Simulated performance of error correction codes in higher dimensions](#), by James Keppens, Quinten Eggerickx, Vukan Levajac, George Simion, and Bart Sorée, arXiv, February 2025 (14 pages).
- [1871] [Transversal STAR architecture for megaquop-scale quantum simulation with neutral atoms](#), by Refaat Ismail, Milan Kornjača, et al., arXiv, September 2025 (32 pages).
- [1872] [Repeated ancilla reuse for logical computation on a neutral atom quantum computer](#), by J. A. Muniz, M. A. Norcia, et al., arXiv, June 2025 (15 pages).
- [1873] [Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays](#), by Yue Wu, Shimon Kolkowitz, Shruti Puri, and Jeff D Thompson, arXiv, January 2022 (16 pages).
- [1874] [Introducing Helios: The Most Accurate Quantum Computer in the World](#), by Quantinuum, Quantinuum news, November 2025.
- [1875] [Experimental validation of a compact fault-tolerant architecture for trapped ions](#), by Noah Berthussen, David Hayes, Tzvetan S. Metodi, Andrew C. Potter, et al., arXiv, September 2026 (18 pages).
- [1876] [Computing with many encoded logical qubits beyond break-even](#), by Shival Dasu, David Hayes, Ben Criger, Andrew C. Potter, David Amaro, et al., arXiv, February 2026 (44 pages).
- [1877] [Distributed fault-tolerant quantum memories over a 2xL array of qubit modules](#), by Edwin Tham, Min Ye, Ilia Khait, John Gamble, and Nicolas Delfosse, arXiv, August 2025 (9 pages).
- [1878] [Optimized Clifford Noise Reduction: Theory, Simulations and Experiments](#), by Edwin Tham and Nicolas Delfosse, Quantum Journal, August 2025 (14 pages).
- [1879] [Break-even demonstration of quantum low-density parity-check codes](#), by Edwin Tham, Michael L. Goldman, Shantanu Debnath, Ashay N. Patel, Jyothi Saraladevi, Jason Nguyen, Erik Nielsen, Neal Pienti, Kenneth Wright, John Gamble, and Nicolas Delfosse, arXiv, June 2026 (16 pages).
- [1880] [Correction of chain losses in trapped ion quantum computers](#), by Nolan J. Coble, Min Ye, and Nicolas Delfosse, arXiv, November 2025 (11 pages).
- [1881] [Holographic quantum codes with trapped ions](#), by Alex Steiner, Gerard Anglès Munné, Robert Freund, Ivan Pogorelov, Michael Meth, Robert J. Harris, Gavin Brennen, Thomas M. Stace, Thomas Monz, Rainer Blatt, Felix Huber, and Martin Ringbauer, arXiv, July 2026 (10 pages).
- [1882] [Experimental Repetitive Quantum Error Correction](#), by Philipp Schindler, Julio T. Barreiro, Thomas Monz, Volckmar Nebendahl, Daniel Nigg, Michael Chwalla, Markus Hennrich, and Rainer Blatt, Science, May 2011 📄.
- [1883] [Experimental Quantum Computations on a Topologically Encoded Qubit](#), by Daniel Nigg, Thomas Monz, Rainer Blatt, et al., Science, June 2014 (20 pages) 📄 (preprint on arXiv).
- [1884] [Entangling logical qubits with lattice surgery](#), by Alexander Erhard, Thomas Monz, et al., Nature, January 2021 📄 (preprint on arXiv).
- [1885] [Experimental deterministic correction of qubit loss](#), by Roman Stricker, Davide Vodola, Alexander Erhard, Lukas Postler, Michael Meth, Martin Ringbauer, Philipp Schindler, Thomas Monz, Markus Müller, and Rainer Blatt, Nature, September 2020 📄 (preprint on arXiv).
- [1886] [Genuine Multipartite Entanglement between Logical Qubits via Cross-Code Lattice Surgery](#), by Alex Steiner, Nicolai Friis, et al., arXiv, July 2026 (13 pages).
- [1887] [Realization of Universal Ion Trap Quantum Computation with Decoherence Free Qubits](#), by Thomas Monz, Rainer Blatt, et al., arXiv, September 2009 (4 pages).
- [1888] [Characterizing physical and logical errors in a transversal CNOT via cycle error reconstruction](#), by Nicholas Fazio, Stephen D. Bartlett, et al., Physical Review X, February 2026 (18 pages).
- [1889] [Scaling and logic in the colour code on a superconducting quantum processor](#), by N. Lacroix, Craig Gidney, K. J. Satzinger, et al., Nature, May 2025 (7 pages).
- [1890] [Extractors: QLDPC Architectures for Efficient Pauli-Based Computation](#), by Zhiyang He, Alexander Cowtan, Dominic J. Williamson, and Theodore J. Yoder, arXiv, October 2025 (48 pages).
- [1891] [IBM lays out clear path to fault-tolerant quantum computing](#), by Ryan Mandelbaum et al., IBM blog, June 2025.
- [1892] [Encoding a magic state with beyond break-even fidelity](#), by Riddhi S. Gupta, Benjamin J. Brown, et al., Nature, January 2024.
- [1893] [Universal fault tolerant quantum computation in 2D without getting tied in knots](#), by Margarita Davydova, Benjamin J. Brown, et al., arXiv, May 2025 (62 pages).

- [1894] [Distilling Magic States in the Bicycle Architecture](#), by Shifan Xu, Kun Liu, Patrick Rall, Zhiyang He, and Yongshan Ding, arXiv, February 2026 (13 pages).
- [1895] [Constant-Overhead Addressable Gates via Single-Shot Code Switching](#), by Louis Golowich, Kathleen Chang, and Guanyu Zhu, arXiv, October 2025 (77 pages).
- [1896] [Quantum error detection in qubit-resonator star architecture](#), by Florian Vigneau, Johannes Heinsoo, et al., arXiv, August 2025 (21 pages).
- [1897] [LDPC-cat codes for low-overhead quantum computing in 2D](#), by Diego Ruiz, Jérémie Guillaud, Anthony Leverrier, Mazyar Mirrahimi, and Christophe Vuillot, arXiv, February 2024 (23 pages).
- [1898] [Elevator Codes: Concatenation for resource-efficient quantum memory under biased noise](#), by Peter Shanahan and Diego Ruiz, arXiv, January 2026 (9 pages).
- [1899] [Alice & Bob and Inria Improve Efficiency of Magic State Preparation to Enable Useful Quantum Computing](#), by Niccolò Coppola, Alice & Bob, August 2025.
- [1900] [Interleaving: Modular architectures for fault-tolerant photonic quantum computing](#), by Hector Bombin, Isaac H. Kim, Daniel Litinski, Naomi Nickerson, Mihir Pant, Fernando Pastawski, Sam Roberts, and Terry Rudolph, arXiv, March 2021 (22 pages).
- [1901] [Fusion-based quantum computation](#), by Sara Bartolucci, Patrick Birchall, Hector Bombin, Hugo Cable, Chris Dawson, Mercedes Gimeno-Segovia, Eric Johnston, Konrad Kieling, Naomi Nickerson, Mihir Pant, Fernando Pastawski, Terry Rudolph, and Chris Sparrow, arXiv, January 2021 (25 pages).
- [1902] [Comparison of schemes for highly loss tolerant photonic fusion based quantum computing](#), by Sara Bartolucci, Andrzej Veitia, et al., arXiv, June 2025 (5 pages).
- [1903] [Blueprint for a Scalable Photonic Fault-Tolerant Quantum Computer](#), by J. Eli Bourassa, Ish Dhand, et al., arXiv, February 2021 (38 pages).
- [1904] [Integrated photonic source of Gottesman–Kitaev–Preskill qubits](#), by M. V. Larsen, D. H. Mahler, et al., Nature, June 2025.
- [1905] [Linear-Optical Quantum Computation with Arbitrary Error-Correcting Codes](#), by Blayne W. Walshe, Physical Review Letters, March 2025 (preprint on arXiv).
- [1906] [A Spin-Optical Quantum Computing Architecture](#), by Grégoire de Glinasty, Paul Hilaire, Pierre-Emmanuel Emeriau, Stephen C. Wein, Alexia Salavrakos, and Shane Mansfield, arXiv, July 2024 (29 pages).
- [1907] [Enhanced Fault-tolerance in Photonic Quantum Computing: Comparing the Honeycomb Floquet Code and the Surface Code in Tailored Architecture](#), by Théo Dessertaine, Boris Bourdoncle, Aurélie Denys, Grégoire de Glinasty, Pierre Colonna d'Istria, Gerard Valentí-Rojas, Shane Mansfield, and Paul Hilaire, arXiv, January 2026 (28 pages).
- [1908] [Nonlinear photonic architecture for fault-tolerant quantum computing](#), by Maike Ostmann, Joshua Nunn, and Alex E. Jones, arXiv, October 2025 (9 pages).
- [1909] [Predicting Magic from Very Few Measurements](#), by J. M. Varela, L. L. Keller, A. de Oliveira Junior, D. A. Moreira, R. Chaves, and R. A. Macêdo, arXiv, February 2026 (20 pages).
- [1910] [Quantum Contextuality](#), by Mladen Pavicic, arXiv, March 2023 (68 pages).
- [1911] [Contextuality supplies the ‘magic’ for quantum computation](#), by Mark Howard, Joel Wallman, Victor Veitch, and Joseph Emerson, Nature, June 2014 (preprint on arXiv).
- [1912] [Universal quantum computation with ideal Clifford gates and noisy ancillas](#), by Sergey Bravyi and Alexei Kitaev, PRA, February 2005 (15 pages) (preprint on arXiv).
- [1913] [A fault-tolerant non-Clifford gate for the surface code in two dimensions](#), by Benjamin J. Brown, May 2020.
- [1914] [Logical Magic State Preparation with Fidelity Beyond the Distillation Threshold on a Superconducting Quantum Processor](#), by Yangsen Ye, Jian-Wei Pan, et al., arXiv, May 2023 (12 pages).
- [1915] [Efficient Magic State Distillation by Zero-Level Distillation](#), by Tomohiro Itogawa, Yugo Takada, Yutaka Hirano, and Keisuke Fujii, PRX Quantum, June 2025 (11 pages).
- [1916] [Orchestrating multi-level magic state distillation: a dynamic pipeline architecture](#), by Junshi Wang and Prakash Murali University of Cambridge, arXiv, September 2025 (16 pages).
- [1917] [Reducing Postselection Overhead in Magic-State Cultivation by In-Patch Multiplexing](#), by Dongmin Kim, Jeonggeun Seo, Aniket Patra, and Youngsun Han, arXiv, May 2026 (21 pages).
- [1918] [Magic state cultivation: growing T states as cheap as CNOT gates](#), by Craig Gidney, Noah Shutty, and Cody Jones, arXiv, September 2024 (33 pages).
- [1919] [Magic State Cultivation](#), by Craig Gidney, Simons Institute Quantum Colloquium, April 2025 (61 slides).
- [1920] [Magic state cultivation on a superconducting quantum processor](#), by Emma Rosenfeld, Craig Gidney, et al., arXiv, December 2025 (28 pages).
- [1921] [Magic State Cultivation on the Surface Code](#), by Yotam Vaknin, Shoham Jacoby, Arne Grimsmo, and Alex Retzker, arXiv, February 2025 (10 pages).
- [1922] [Fold-transversal surface code cultivation](#), by Kaavya Sahay, Shruti Puri, et al., arXiv, September 2025 (18 pages).
- [1923] [Efficient magic state cultivation with lattice surgery](#), by Yutaka Hirano, Riki Toshio, Tomohiro Itogawa, and Keisuke Fujii, arXiv, October 2025 (12 pages).
- [1924] [Efficient Magic State Cultivation on \$\mathbb{RP}^2\$](#) , by Zi-Han Chen, Jian-Wei Pan, et al., arXiv, March 2025 (15 pages).
- [1925] [MagicPool: Dealing with Magic State Distillation Failures on Large-Scale Fault-Tolerant Quantum Computer](#), by Yutaka Hirano, Yasunari Suzuki, and Keisuke Fujii, arXiv, July 2024 (13 pages).
- [1926] [A lower bound on the overhead of quantum error correction in low dimensions](#), by Nouédy Baspin, Omar Fawzi, and Ala Shayeghi, arXiv, February 2023 (21 pages).
- [1927] [Borrowed Identities: Malleable Distillation Factories and a Unified Numerical Search](#), by Shraddha Singh, Craig Gidney, and Cody Jones, arXiv, 2026 (11 pages).
- [1928] [A Resource Comparison of Logical T-State Preparation](#), by Jianshuo Gao, Xiao Yuan, and Yuan Yao, arXiv, May 2026 (15 pages).
- [1929] [Efficient Magic State Cultivation for \$\sqrt{T}\$ Gates](#), by I-Chi Chen, Matheus da Silva Fonseca, and Andrew Sornborger, arXiv, June 2026 (11 pages).
- [1930] [Magic tricycles: efficient magic state generation with finite block-length quantum LDPC codes](#), by Varun Menon, J. Pablo Bonilla-Ataides, Mikhail D. Lukin, et al., arXiv, August 2025 (63 pages).
- [1931] [Scheduling Lattice Surgery with Magic State Cultivation](#), by Steven Hofmeyr, Costin Iancu, et al., arXiv, December 2025 (11 pages).
- [1932] [Bridging Superconducting and Neutral-Atom Platforms for Efficient Fault-Tolerant Quantum Architectures](#), by Xiang Fang, Jixuan Ruan, Sharanya Prabhu, Ang Li, Travis Humble, Dean Tullsen, and Yufei Ding, arXiv, January 2026 (15 pages).
- [1933] [Heterogeneous architectures enable a \$138\times\$ reduction in physical qubit requirements for fault-tolerant quantum computing under detailed accounting](#), by Pranav S. Mundada, Aleksei Khindanov, Yulun Wang, Claire L. Edmunds, Paul Coote, Michael J. Biercuk, Yuval Baum, and Michael Hush, arXiv, April 2026 (34 pages).
- [1934] [Quantum error correction and fault tolerance: A comprehensive tutorial](#), by Daniel J. Spencer, Shubham P. Jain,

- Andrew Tanggara, Zeen Sun, Tobias Haug, Derek Khu, and Kishor Bharti, arXiv, May 2026 (209 pages).
- [1935] [A comparative code study for quantum fault-tolerance](#), by Andrew W. Cross, David P. DiVincenzo, and Barbara M. Terhal, arXiv, February 2009 (33 pages).
- [1936] [Quantum codes do not increase fidelity against isotropic errors](#), by J. Lacalle, L.M. Pozo-Coronado, A.L. Fonseca de Oliveira, and R. Martin-Cuevas, arXiv, January 2022 (18 pages).
- [1937] [Quantum computing is scalable on a planar array of qubits with fabrication defects](#), by Armands Strikis, Simon C. Benjamin, and Benjamin J. Brown, arXiv, June 2023 (18 pages).
- [1938] [Midcircuit Operations Using the omg Architecture in Neutral Atom Arrays](#), by M. A. Norcia et al., Physical Review X, November 2023.
- [1939] [Demonstration of the trapped-ion quantum CCD computer architecture](#), by J. M. Pino, B. Neyenhuis, et al., Nature, April 2021 (preprint on arXiv).
- [1940] [Measurement-Induced State Transitions in a Superconducting Qubit: Beyond the Rotating Wave Approximation](#), by Daniel Sank, John M. Martinis, et al., Physical Review Letters, 2016 (preprint on arXiv).
- [1941] [Microarchitectures for Heterogeneous Superconducting Quantum Computers](#), by Samuel Stein, Andrew A. Houck, Isaac L. Chuang, Michael Austin DeMarco, et al., arXiv, May 2023 (16 pages).
- [1942] [Do Not Let CNOTs Overwhelm the Decoder: Scheduling Transversal Gates for Fast FTQC](#), by Shota Ikari, Yuga Hirai, Yasunari Suzuki, Hiroshi Nakamura, and Yosuke Ueno, arXiv, August 2026 (16 pages).
- [1943] [A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery](#), by Daniel Litinski, arXiv, February 2019 (29 pages).
- [1944] [A Game of Surface Codes: Large-Scale Quantum Comp. w. Lattice Surgery](#), by Daniel Litinski (FU Berlin), QEC'19, August 2019 (49 mn).
- [1945] [Applications of universal parity quantum computation](#), by Michael Fellner, PRA, October 2022 (8 pages).
- [1946] [Universal Parity Quantum Computing](#), by Michael Fellner, PRL, October 2022 (7 pages).
- [1947] [Error-corrected Hadamard gate simulated at the circuit level](#), by György P. Gehér, Campbell McLauchlan, Earl T. Campbell, Alexandra E. Moylett, and Ophelia Crawford, arXiv, June 2024 (26 pages).
- [1948] [Exploiting Movable Logical Qubits for Lattice Surgery Compilation](#), by Laura S. Herzog, Lucas Berent, Aleksander Kubica, and Robert Wille, arXiv, December 2025 (7 pages).
- [1949] [Assessing requirements to scale to practical quantum advantage](#), by Michael E. Beverland, Alexander Vashchillo, et al., arXiv, November 2022 (41 pages).
- [1950] [Low overhead quantum computation using lattice surgery](#), by Austin G. Fowler and Craig Gidney, arXiv, August 2019 (15 pages).
- [1951] [Lattice Surgery for Dummies](#), by Avimita Chatterjee, Subrata Das, and Swaroop Ghosh, arXiv, April 2024 (10 pages).
- [1952] [New magic state distillation factories optimized by temporally encoded lattice surgery](#), by Prithviraj Prabhu and Christopher Chamberland, arXiv, October 2022 (35 pages).
- [1953] [A SAT Scalpel for Lattice Surgery: Representation and Synthesis of Subroutines for Surface-Code Fault-Tolerant Quantum Computing](#), by Daniel Bochen Tan, Craig Gidney, et al., arXiv, August 2024 (15 pages).
- [1954] [Spatially parallel decoding for multi-qubit lattice surgery](#), by Sophia Fuhui Lin, Eric C. Peterson, Krishanu Sankar, and Prasanth Sivarajah, arXiv, March 2025 (13 pages).
- [1955] [Towards Lattice Surgery Compilation for the Color Code Using Pipe Diagrams](#), by Laura S. Herzog, Gilad Kishony, Robert Wille, and Austin Fowler, arXiv, July 2026 (20 pages).
- [1956] [Optimization by Decoded Quantum Interferometry](#), by Stephen P. Jordan, Ryan Babbush, et al., Nature, October 2025 (7 pages).
- [1957] [The FLuid Allocation of Surface code Qubits \(FLASQ\) cost model for early fault-tolerant quantum algorithms](#), by William J. Huggins, Tanuj Khattar, Amanda Xu, Matthew Harrigan, Christopher Kang, Guang Hao Low, Austin Fowler, Nicholas C. Rubin, and Ryan Babbush, arXiv, November 2025 (60 pages).
- [1958] [How to factor 2048 bit RSA integers with less than a million noisy qubits](#), by Craig Gidney, arXiv, May 2025 (40 pages).
- [1959] [Efficient fault-tolerant code switching via one-way transversal CNOT gates](#), by Sascha Heußen and Janine Hilder, arXiv, September 2024 (27 pages).
- [1960] [Experimental fault-tolerant code switching](#), by Ivan Pogorelov, Markus Müller, and Thomas Monz, arXiv, March 2024 (18 pages).
- [1961] [Code switching revisited: Low-overhead magic state preparation using color codes](#), by Lucas Daguerre and Isaac H. Kim, arXiv, April 2025 (12 pages).
- [1962] [Optimization of Quantum Error Correcting Code under Temporal Variation of Qubit Quality](#), by Subrata Das and Swaroop Ghosh, arXiv, May 2025 (6 pages).
- [1963] [Decoder Switching: Breaking the Speed-Accuracy Tradeoff in Real-Time Quantum Error Correction](#), by Riki Toshio, Keisuke Fujii, et al., arXiv, October 2025 (26 pages).
- [1964] [Fault-tolerant interface between quantum memories and quantum processors](#), by Hendrik Poulsen Nautrup, Nicolai Friis, and Hans J. Briegel, arXiv, August 2017 (11 pages).
- [1965] [Decoding across Transversal Clifford Gates in the Surface Code](#), by Marc Serra-Peralta, PRX Quantum, February 2026 (41 pages).
- [1966] [Fault-Tolerant Error Detection Above Break-Even for Multi-Qubit Gates](#), by Colburn Riffel, Reece Robertson, and Peter Hendrickson, arXiv, April 2026 (7 pages).
- [1967] [No Tile Left Behind: Multiprogramming for Surface-Code Architectures](#), by Archisman Ghosh, Avimita Chatterjee, and Swaroop Ghosh, arXiv, April 2026 (11 pages).
- [1968] [Toward designing workload-aware Surface Code Architectures](#), by Archisman Ghosh, Avimita Chatterjee, and Swaroop Ghosh, arXiv, April 2026 (14 pages).
- [1969] [Universal quantum computation with group surface codes](#), by Naren Manjunath, Vieri Mattei, Apoorv Tiwari, and Tyler D. Ellison, arXiv, March 2026 (55 pages).
- [1970] [Spacetime-Efficient and Hardware-Compatible Complex Quantum Logic Units in qLDPC Codes](#), by Willers Yang, Jason Chadwick, Mariesa H. Teo, Joshua Vizslai, and Fred Chong, arXiv, March 2026 (26 pages).
- [1971] [Bounded-depth spacetime lattice surgery for resource-efficient fault-tolerant quantum computation](#), by Kou Hamada, Hiroki Hamaguchi, Yosuke Ueno, Yasunari Suzuki, Teruo Tanimoto, and Nobuyuki Yoshioka, arXiv, June 2026 (31 pages).
- [1972] [Factoring 2048-bit RSA Integers in 177 Days with 13436 Qubits and a Multimode Memory](#), by Élie Gouzien and Nicolas Sangouard, arXiv, September 2021 (5 pages).
- [1973] [LSQCA: Resource-Efficient Load/Store Architecture for Limited-Scale Fault-Tolerant Quantum Computing](#), by Takumi Kobori, Yasunari Suzuki, Yosuke Ueno, Teruo Tanimoto, Syngé Todo, and Yuuki Tokunaga, arXiv, April 2025 (17 pages).
- [1974] [Continuous-Variable Fault-Tolerant Quantum Computation under General Noise](#), by Takaya Matsuura, Nicolas C. Menicucci, and Hayata Yamasaki, arXiv, October 2024 (12 pages).
- [1975] [Towards fault-tolerant quantum computation with universal continuous-variable gates](#), by Sheron Blair, Francesco Arzani, Giulia Ferrini, and Alessandro Ferraro, arXiv, June 2025 (10 pages).

- [1976] [Synchronization for Fault-Tolerant Quantum Computers](#), by Satvik Maurya and Swamit Tannu, ISCA'25 Synchronization for Fault-Tolerant Quantum Computers, June 2025 (16 pages) 📄 (preprint on [arXiv](#)).
- [1977] [A Three-Layer Architecture for Fault-Tolerant Quantum Computing](#), by Zhirao Wang, Zhou You, Yiming Huang, Tianyi Li, Ying Li, Xiao Yuan, and Yuan Yao, [arXiv](#), June 2026 (27 pages).
- [1978] [Fault Tolerant Non-Clifford State Preparation for Arbitrary Rotations](#), by Hyeonrak Choi, Dirk Englund, Yongshan Ding, et al., [arXiv](#), March 2023 (14 pages).
- [1979] [Quantum algorithms for solving a drift-diffusion equation: analysing circuit depths](#), by Ellen Devereux and Animesh Datta, [arXiv](#), September 2025 (10 pages).
- [1980] [Practical quantum advantage on partially fault-tolerant quantum computer](#), by Riki Toshio, Keisuke Fujii, et al., [arXiv](#), August 2024 (40 pages).
- [1981] [STAR-Magic Mutation: Even More Efficient Analog Rotation Gates for Early Fault-Tolerant Quantum Computer](#), by Riki Toshio, Shota Kanasugi, Jun Fujisaki, Hirotaka Oshima, Shintaro Sato, and Keisuke Fujii, [arXiv](#), March 2026 (26 pages).
- [1982] [General circuit compilation protocol into partially fault-tolerant quantum computing architecture](#), by Tomochika Kurita, [arXiv](#), March 2026 (11 pages).
- [1983] [Partially Fault-Tolerant Quantum Computation for Megaquop Applications](#), by Ming-Zhi Chung, Alan Ho, Masoud Mohseni, Pooya Ronagh, John Martinis, et al., [arXiv](#), March 2026 (30 pages).
- [1984] [Fast and Parallel High-Rate STAR Architecture for Megaquop Quantum Simulation](#), by Refaat Ismail, Milan Kornjača, Hong-Ye Hu, Nishad Maskara, Sheng-Tao Wang, Hengyun Zhou, and Chen Zhao, [arXiv](#), June 2026 (29 pages).
- [1985] [FT-Weave: Real-Time Compilation Framework for Reconfigurable Fault-Tolerant Quantum Architectures](#), by Wan-Hsuan Lin, Milan Kornjača, Chen Zhao, Sheng-Tao Wang, and Jason Cong, [arXiv](#), September 2026 (28 pages).
- [1986] [RESCQ: Realtime Scheduling for Continuous Angle Quantum Error Correction Architectures](#), by Sayam Sethi and Jonathan Mark Baker, [arXiv](#), March 2025 (16 pages).
- [1987] [Fujitsu and The University of Osaka develop new technologies for chemical material energy calculations on early-FTQC quantum computers](#), by Fujitsu and the University of Osaka, March 2026 (28 slides).
- [1988] [Error-structure-tailored early fault-tolerant quantum computing](#), by Pei Zeng, Guo Zheng, Qian Xu, and Liang Jiang, [arXiv](#), November 2025 (35 pages).
- [1989] [Suppressing errors in analog logical rotation gates via balanced fusion](#), by Sam McArdle, Alexander M. Dalzell, and Fernando G.S.L. Brandão, [arXiv](#), July 2026 (17 pages).
- [1990] [Fire and ice: Partially fault-tolerant quantum computing with selective state filtering](#), by Ben W. Reichardt, David Aasen, and Rui Chao, [arXiv](#), May 2026 (14 pages).
- [1991] [Theory of quasi-exact fault-tolerant quantum computing and valence-bond-solid codes](#), by Dong-Sheng Wang, Raymond Laflamme, et al., [arXiv](#), May 2021 (22 pages).
- [1992] [Complexity and order in approximate quantum error-correcting codes](#), by Jinmin Yi, Weicheng Ye, Daniel Gottesman, and Zi-Wen Liu, [Nature Physics](#), September 2024 📄 (preprint on [arXiv](#)).
- [1993] [Approximate Quantum Error Correction with 1D Log-Depth Circuits](#), by Guoding Liu et al., [PRX Quantum](#), February 2026 (76 pages).
- [1994] [Complexity and order in approximate quantum error-correcting codes](#), by Jinmin Yi, Weicheng Ye, Daniel Gottesman, and Zi-Wen Liu, [arXiv](#), September 2024 (30 pages).
- [1995] [NISQ+: Boosting quantum computing power by approximating quantum error correction](#), by Adam Holmes, Fred-eric T. Chong, et al., [arXiv](#), April 2020 (13 pages).
- [1996] [Quantum Computing in the NISQ era and beyond](#), by John Preskill, [Quantum Journal](#), July 2018 (20 pages).
- [1997] [QIP 2022 |Software of QIP, by QIP, and for QIP](#), by Dave Bacon from Google, March 2022.
- [1998] [Beyond NISQ: The Megaquop Machine](#), by John Preskill, [Quantum Frontiers](#), December 2024 (9 pages) 📄 (preprint on [arXiv](#)).
- [1999] [Early Fault-Tolerant Quantum Computing](#), by Amara Katabarwa et al., [PRX Quantum](#), November 2023.
- [2000] [On proving the robustness of algorithms for early fault-tolerant quantum computers](#), by Rutuja Kshirsagar, Amara Katabarwa, and Peter D. Johnson, [arXiv](#), February 2025 (27 pages).
- [2001] [Modeling the Performance of Early Fault-Tolerant Quantum Algorithms](#), by Qiyao Liang et al., [Physical Review Research](#), December 2023.
- [2002] [Variational Quantum Algorithms in the era of Early Fault Tolerance](#), by Siddharth Dangwal, Gokul Subramanian Ravi, et al., [arXiv](#), May 2025 (15 pages).
- [2003] [Contrasting Statistical Phase Estimation with the Variational Quantum Eigensolver in the era of Early Fault Tolerant Quantum Computation](#), by Ming-Zhi Chung, Ryosuke Imai, et al., [arXiv](#), September 2024 (30 pages).
- [2004] [Robust sparse IQP sampling in constant depth](#), by Louis Paletta, Anthony Leverrier, Alain Sarlette, Maziar Mirrahimi, and Christophe Vuillot, [Quantum Journal](#), July 2023 (16 pages).
- [2005] [Mind the gaps: The fraught road to quantum advantage](#), by Jens Eisert and John Preskill, [arXiv](#), October 2025 (11 pages).
- [2006] [Quantum Computing – A new scientific revolution in the making](#), by Koen Bertels, Imran Ashraf, et al., [arXiv](#), May 2024 (16 pages).
- [2007] [Upper Bounds for the Clock Speeds of Fault-Tolerant Distributed Quantum Computation using Satellites to Supply Entangled Photon Pairs](#), by Hudson Leone, Simon Devitt, et al., [arXiv](#), February 2024 (11 pages).
- [2008] [Resource Estimation for Fault-Tolerant Quantum Programs](#), by Bonan Su, Yuan Feng, Li Zhou, and Mingsheng Ying, [arXiv](#), August 2026 (29 pages).
- [2009] [An iterative transversal CNOT decoder](#), by Kwok Ho Wan, Mark Webber, Austin G. Fowler, and Winfried K. Hensinger, [arXiv](#), April 2025 (21 pages).
- [2010] [Engineering CSS surgery: compiling any CNOT in any code](#), by Clément Poirson, Joschka Roffe, and Robert I. Booth, [arXiv](#), June 2025 (49 pages).
- [2011] [Lattice surgery realized on two distance-three repetition codes with superconducting qubits](#), by Ilya Besedin, Markus Müller, Andreas Wallraff, et al., [Nature Physics](#), January 2026 (9 pages).
- [2012] [Lattice surgery for near-term experimental logical qubit entanglement creation in planar architectures](#), by Lukas Bödeker, Áron Márton, Luis Colmenarez, Ilya Besedin, Andreas Wallraff, and Markus Müller, [arXiv](#), June 2026 (20 pages).
- [2013] [When T-Depth Misleads: Predicting Fault-Tolerant Quantum Execution Slowdown under Magic-State Delivery Constraints](#), by Boshuai Ye, Arif Ali Khan, and Peng Liang, [arXiv](#), May 2026 (10 pages).
- [2014] [Price and Payoff: Non-Determinism in Fault Tolerant Quantum Computation](#), by Aditi Awasthi, Sayam Sethi, Sahil Khan, Gokul Subramanian Ravi, and Jonathan Mark Baker, [arXiv](#), May 2026 (11 pages).
- [2015] [Scalable Constant-Time Logical Gates for Large-Scale Quantum Computation Using Window-Based Correlated Decoding](#), by Jiaxuan Zhang, Guo-Ping Guo, et al., [arXiv](#), January 2025 (12 pages).
- [2016] [Design Fault Tolerant Quantum Computing applications with the open-source Resource Estimator](#), by Fabrice Frachon et al., January 2024.

- [2017] [Expressing and Analyzing Quantum Algorithms with Qualtran](#), by Matthew P. Harrigan, Ryan Babbush, Nicholas C. Rubin, et al., arXiv, September 2024 (39 pages).
- [2018] [PsiQuantum announces QREF and Bartiq: Open-Source Software for Better Tools, Libraries and Datasets](#), by PsiQuantum, May 2024.
- [2019] [StreamingQEC: Streaming Quantum Error Correction in Tightly Integrated Quantum-Classical Systems via Certified Recurrence](#), by Panayiotis Christou, Shuwen Kan, Hao Wang, and Ying Mao, arXiv, July 2026 (49 pages).
- [2020] [Factoring 2048 bit RSA integers with a half-million-qubit modular atomic processor](#), by Tian Xue and Jacob P. Covey, arXiv, May 2026 (29 pages).
- [2021] [How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits](#), by Craig Gidney and Martin Ekerå, arXiv, April 2021 (31 pages).
- [2022] [Shor's algorithm is possible with as few as 10,000 reconfigurable atomic qubits](#), by Madelyn Cain, Qian Xu, Robbie King, Lewis R. B. Picard, Harry Levine, Manuel Endres, John Preskill, Hsin-Yuan Huang, and Dolev Bluvstein, arXiv, March 2026 (15 pages).
- [2023] [Resource Analysis of Low-Overhead Transversal Architectures for Reconfigurable Atom Arrays | Proceedings of the 52nd Annual International Symposium on Computer Architecture](#), by Hengyun Zhou, Dolev Bluvstein, Mikhail D. Lukin, et al., ACM Conferences, June 2025 (17 pages)  (preprint on [arXiv](#)).
- [2024] [Performance Analysis of a Repetition Cat Code Architecture: Computing 256-bit Elliptic Curve Logarithm in 9 Hours with 126133 Cat Qubits](#), by Élie Gouzien, Diego Ruiz, Francois-Marie Le Régent, Jérémie Guillaud, and Nicolas Sangouard, arXiv, August 2023 (37 pages).
- [2025] [Computing 256-bit elliptic curve discrete logarithms in 26 days on a fault-tolerant trapped-ion quantum computer with 20,000 qubits](#), by Thomas Häner, Nicolas Delfosse, John Gamble, Martin Roetteler, et al., arXiv, September 2026 (70 pages).
- [2026] [The Grand Challenge of Quantum Applications](#), by Ryan Babbush, Jarrod R. McClean, Thomas O'Brien, Nicholas C. Rubin, et al., arXiv, November 2025 (22 pages).
- [2027] [Strategic Plan for Neutral Atom Quantum Computation](#), by Adrian J. Menssen, Antoine Browaeys, Dirk Englund, Liang Jiang, Thierry Lahaye, Mikhail D. Lukin, Mark Saffman, Pascal Scholl, Vladan Vuletic, Hengyun Zhou, et al., arXiv, July 2026 (135 pages).

This **Error handling** subsection contains 706 bibliographical references and notes containing at least 16,755 pages.

5.3 Quantum memory

Quantum computers implement something that resembles what is called “in-memory processing” in classical computing, an architecture where operations are applied directly in memory. A quantum processor holds its quantum state in its qubits while gates act on that state, and readout happens at the end with quantum measurement. But a quantum register is not a memory architecture in the computer-architecture sense. It is not classically addressable and its vector state representing its internal content cannot be read directly.

5.3.1 Classical memory

This model is not the one currently used in classical computing where memory and data storage are located in a different place than processors, whether they are CPU or GPUs. Based on the John von Neumann model invented in the mid 1940s, the processor is constantly pulling data from memory (usually, a random access memory, RAM) and putting it in the CPU cache and register. Operations are made using a very small number of bytes stored in registers and then sent back on the RAM. Classical memory is organized in a complex hierarchy with, viewed from the processor contains: registers, several layers of cache, RAM, and then storage itself potentially decomposed between short term storage (using SSD drives) and longer term storage (hard disks and even magnetic tapes that are still around) (Figure 288).




Figure 288: The memory hierarchy in classical computers. The larger, the slower! The data flow is a bit different, the CPU or specialized input/output units within the CPU being dedicated to data transfer between these different layers of memory.

There is even a technology distinction between the cache RAM that sits inside the CPU and GPU and that is based on SRAM technology, which is fast but energetically costly, and the outside memory sitting in separate chips, using DRAM, which is slower, has a larger capacity, but is less expensive (Figure 289).

All these classical memories require error correction. Typically, ECC (Error Correction Code) bits are added with 8 additional bits for every 64-bit word in the usual SEC-DED scheme, 7 bits being enough only for single-error correction without double-error detection, generating a 11 to 13% area overhead. Level 1 caches often settle for simple parity, while L2, L3 and external DRAM use SEC-DED with this same 72/64 ratio [2028].

memory type	SRAM	DRAM
power supply	constant	periodical refresh
power consumption ratio	3 to 5	1
transistors per storage bit	6	1
cost ratio	100	1
capacity	low (<100 MB in Nvidia GPUs)	high (in GBs)
access time ratio	3 – 5	1
use cases	CPU/GPU cache	primary memory

Figure 289: Key differences between SRAM and DRAM memory. SRAM is expensive and very fast while DRAM is slower but cheaper. So, SRAM is used as cache memory within processors while DRAM is used as external primary memory, like in the HBM3E memories used along GPUs. © Olivier Ezratty, August 2025.

Are there in-memory processing architectures in the classical computing realm? Yes, but they are not commonplace. It’s the case for neuromorphic processors where artificial neurons are connected with each other to solve some neural network problems. Memories exist that embed small processing cores, like the ones from UPMEM (France), whose DRAM Processing Units are proprietary 32-bit in-order RISC cores running at 350 to 500 MHz with up to 24 hardware threads, neither Arm nor RISC-V, each paired with a 64 MB DRAM bank.

In the classical computing world, the current focus is to make data transfer faster between the different layers in the memory stack, like with the HBM memory that is next to the Nvidia GPUs. It currently operates at an input/output data speed of 8 TB/s per GPU on Nvidia Blackwell B200 and B300 parts, against 4.8 TB/s on the previous H200 generation.

This is something to keep in mind when comparing classical and quantum computers, these being much slower to ingest data due to the low gate speeds that, at best are at 10 ns, and the cost of error correction that slows down everything.

5.3.2 Quantum algorithms requirements

One could view the need for quantum memories as similar to what happens in classical computing, with architectures using different layers of quantum memories to optimize data flows. The QPU qubit register is a fast but costly memory dedicated to processing. Additional

quantum memories could be the equivalent of classical DRAM and serve as buffer memory, particularly for loading and accessing classical data efficiently in the background.

So, “non-processing” quantum memory is some quantum memory capable of storing the quantum state of qubits and then using them to feed quantum computer registers [2029]. It should be able to store superposed and entangled states and deliver it to whatever computing is needed and particularly for oracle based quantum algorithms. It is part of a broader category usually labeled “quantum RAM” or qRAM, a label which covers three different objects that should be kept apart: quantum-accessible classical memory (QACM), where classical data is queried with superposed addresses, coherent quantum memory (QAQM), where the stored content is itself a quantum state, and amplitude loading circuits, which prepare a state whose amplitudes encode a classical vector [2030]. Their costs, their error models and their feasibility are not the same.

One anticipated usage of quantum memory is to temporarily store the state of a qubit register during a data preparation process, a usual lengthy process, before transferring it to a faster quantum processing unit. With N qubits, this memory would be able to store in theory 2^N different computational vector state amplitude values, the exponent being essential here since the whole scaling argument rests on it.

According to the no-cloning theorem, the content of this memory cannot be the copy of the state of other quantum registers. In quantum computing, a quantum memory is used to store data to be later used in quantum processing. Data preparation and encoding depends on the algorithm. It is necessary for certain types of quantum algorithms such as Grover’s search and quantum machine learning algorithms [2035, 2036]. The most demanding encoding is when you encode a vector of 2^N values (well, minus 1 for normalization constraints) in the whole computational state vector [2037].

The normalization constraint and the encoded state vector are:

$$\sum_{i=0}^{2^N-1} |x_i|^2 = 1 \quad \sum_{i=0}^{2^N-1} x_i |i\rangle = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{2^N-1} \end{bmatrix}$$

This creates a superposition with all or some of the basis states from the computational basis. Namely, we encode a vector $\mathbf{x}$ containing 2^N real (or even complex) number values from $\mathbf{x}_0$ to $\mathbf{x}_{2^N-1}$

with the normalization constraint that the sum of the squared moduli of these values is equal to 1.

It ends up creating the state vector on the right with 2^N amplitudes $\mathbf{x}_i$ associated with the vectors $|i\rangle$ from the

computational basis. This is called amplitude encoding. Since this data encoding grows exponentially with the number of qubits, it may erase any computing speedup we would gain later. So, this is efficient only if we find a way to make this fast. One solution is to encode only sparse vectors where only a few values are nonzero.

Quantum data in a quantum memory could potentially be transferred to some fast computing qubits, using variants of quantum state teleportation. Even using quantum error correction, this could be potentially very costly and is not envisioned in most architectures.

5.3.3 Quantum memory types

There are several types of quantum memories used in quantum computing (Figure 290).

qRAM

qRAM was proposed by Vittorio Giovannetti, Seth Lloyd and Lorenzo Maccone in 2008 with the *bucket brigade* addressing system, based on a qutrit tree (three states quantum objects) containing wait/left/right flags [2029, 2038], sort of decision trees to reach the right memory cell. It is also named QACM for Quantum Access Classical Memory [2039] and implicit qRAM.

This quantum addressing system could be used for accessing both classical bits and coherent states in qubits. The first case may be useful when building an oracle function for quantum circuits like a Grover search [2040, 2041]. In the full quantum case, the addresses are in coherent superposition and the memory returns, in superposition, the content stored at each address. What the oracle exposes is the data indexed by an address, classical bits or words for a QACM and a stored quantum state for a coherent memory. It does not give access to an amplitude α_i of an arbitrary quantum state living elsewhere, which is a different operation and is forbidden as such [2031]. The α_j coefficients are the weighted superposition of addresses corresponding to computational basis states $|j\rangle$. One caveat is that the tree size scales badly, but hopefully some workaround this problem are proposed [2042, 2043].

$$\sum_j \alpha_j |j\rangle |0\rangle \rightarrow \sum_j \alpha_j |j\rangle |b_j\rangle$$

This is the result of the query, with weights applied to $|b_j\rangle$ j-th memory location.

This is the best case to obtain some exponential computing time separation for some machine learning algorithms between using or not using some quantum memory [2044]. Throughout this discussion, N denotes the number of memory cells and $n = \log_2 N$ the address width, and the quantities being compared are stated explicitly, since total hardware, active elements per query, circuit depth, address lines and query error




Figure 290: Various classes of quantum memories and use cases. © Olivier Ezratty, 2021-August 2025. Image sources: [2030–2033]. See also this 2026 review: [2034].

are five distinct resources. In classical RAM, an array of $N = 2^n$ cells is usually organized in a 2-dimensional lattice addressed by $O(\sqrt{N})$ row and column lines. What Giovannetti, Lloyd and Maccone compare is not the number of address lines but the number of switches that must be activated per memory call, which is $O(N)$ in a conventional fanout architecture.

In bucket brigade qRAM, this switching cost decreases to $O(\log N)$ to address a particular computational basis vector amplitude. But this must take into account the burden of any quantum error correction for each qRAM router node [2045–2050] and other physics limitations [2051]. Various implementations of the bucket brigade solution have been proposed so far [2052], including one using quantum walks, with the benefit of being more robust to decoherence and easier to parallelize [2053]. Before any qRAM data transfer to computation qubits can be done, an uncompute processing must be implemented that removes the selected computational basis vectors addresses from the related data.

The bucket-brigade physical device has $O(N)$ quantum router nodes. For $N = 2^{50}$ data items (relevant for some HHL applications), this means $\sim 10^{15}$ physical routers. In a naive fanout architecture where every node is actively entangled with the address register, the expected number of faults over $O(N)$ nodes scales as $O(N \cdot \varepsilon)$: with $N = 10^{15}$ and $\varepsilon = 10^{-3}$, one would expect $\sim 10^{12}$ faults per query, meaning a query fidelity indistinguishable from zero.

The bucket brigade avoids exactly this. Because only $O(\log N)$ routers are in the *wait* state along the queried branch at any instant, and because errors occurring on unqueried branches remain confined to those branches, the query infidelity scales polylogarithmically rather than linearly with N . This was first argued from restricted noise models and then proven for arbitrary error

channels, including depolarizing and coherent errors, by Connor Hann et al. in 2021 [2054]. Shen et al. verified it experimentally in 2025 by injecting tunable depolarizing errors into individual routers of a three-layer device: the fidelity slope was -0.255 for the queried branch against -0.005 for the most distant unqueried branch, a fifty-fold difference, with the entanglement entropy of the routers decreasing with tree depth [2055]. Under the noise model analyzed by Hann et al., single-query infidelity therefore scales polylogarithmically rather than linearly with N , which means that a single query does not by itself require every router node to be error corrected. Whether fault-tolerant protection is needed in practice depends on the target error, the number of queries and the router implementation, and the same authors also analyze error-corrected implementations.

The real cost objection lies elsewhere. First, an N -address memory needs $O(N)$ physical routers, which is exponential in the address width: $N = 2^{20}$ is already one million routers for a one-megabit memory. Second, the polylogarithmic infidelity is a *per-query* result. Algorithms consume 10^9 to 10^{15} queries, so the accumulated infidelity forces error correction anyway once the query count is large, which brings back logical-qubit encoding of the routers and turns the memory into a second fault-tolerant system alongside the processor [2048, 2056]. Third, the router qubits must idle coherently for the whole query duration, which is a coherence-time constraint, not a gate-error one. A resource estimate that charges $O(\log N)$ for a query while leaving the $O(N)$ hardware out of the accounting is therefore incomplete, a point made early on by Scott Aaronson [2057].

For encoding a unitary U which is an $n \times n$ matrix, the matrix elements U_{ij} are stored classically at the leaves. To implement the oracle $O_U: |i\rangle |j\rangle |0\rangle \rightarrow |i\rangle |j\rangle |U_{ij}\rangle$, the address register is $|i, j\rangle$ (a $2 \log n$ bit address), and each QRAM query fetches the corresponding matrix

element in superposition. This oracle then feeds into a block-encoding construction. For an s -sparse matrix, one also needs a sparse oracle O_{col} that, given $|i\rangle|k\rangle$, returns the column index of the k -th non-zero element in row i . These two oracles help build the block-encoding circuit:

$$(|0\rangle \otimes I) \cdot U_{\text{BE}} \cdot (|0\rangle \otimes I) = \frac{A}{s \cdot \|A\|_{\max}}$$

via the standard sparse block-encoding circuit [2058, 2059]. The block-encoding unitary U_{BE} calls O_U and O_{col} a constant number of times.

For LCU (Linear Combination of Unitaries), if $A = \sum_k \lambda_k U_k$, two distinct subroutines are involved. PREPARE loads the coefficients λ_k into an ancilla register, which is a data loading task a QROM or a qRAM can perform. SELECT applies the k -th unitary conditioned on the address, which is a structured circuit and not a memory operation: storing a classical description of U_k does not execute it coherently, and its cost depends on how the U_k are represented and synthesized. The SELECT oracle

$$\text{SELECT } |k\rangle|\psi\rangle = |k\rangle U_k |\psi\rangle$$

must coherently execute the k -th unitary given the address k in superposition. This requires not just data retrieval but coherent circuit execution, a “quantum-addressable circuit bank” that is significantly harder to realize than classical data QRAM.

The algorithms claiming exponential or polynomial speedups via QRAM are implicitly assuming the use of a quantum computing device that simultaneously requires more quantum hardware than the computation it enables and has no known error-corrected physical architecture at scale.

qRAM can be extended to manage a quantum database, defined as a balanced superposition state storing data correlated to specific orthogonal index states in a quantum superposition with essential database operations (insert, retrieve, delete, indexing) [2060].

Since qRAM is currently very difficult to practically realize experimentally, researchers are also working on creating classical qRAM simulators, particularly for simulating the effects of noise on memory [2061, 2062].

There are two less common other variants of qRAM:

- **Explicit qRAM** encodes data in physical qubits and then, use quantum circuits to extract the encoded data [2063]. There is no specific addressing system to selectively access parts of this memory. This is the scenario depicted above. Also named QAQM for Quantum Access Quantum Memory [2039, 2064]. There, data is quantum and data access is quantum.

- **Flip-flop qRAM** is a variant of explicit qRAM based on qubits circuit algorithms used to efficiently load classical data in a qubit register [2031, 2065].

QROM

Quantum Read Only Memories allow the retrieval of *classical* data indexed by an address register held in superposition [2066]. A QROM is used to encode an oracle as a quantum circuit, particularly with QPE and HHL.

From a hardware architecture perspective, a QROM is not a physical component. Unlike classical computing, where ROM is a distinct physical memory chip like an EEPROM module on a motherboard, a QROM is a software construct compiled directly into the quantum circuit. It explains why it is not included in Figure 290. The QROM circuit being reused multiple times in the algorithm circuit, some of its elements can be precompiled once like its logical optimization, its space-time layout on the quantum processor and lattice surgery implementation if any.

These figures belong to specific setups. For N data items each of d bits, the reference construction is unary iteration, which costs $4N - 4$ T gates and $O(N)$ depth while needing only $O(\log N)$ ancilla qubits. Its T-count is independent of the word width d . Trading space for T gates with the SelectSwap or QROAM construction, which duplicates the output register λ times, brings the T-count down to $O(N/\lambda + \lambda d)$, minimized at $O(\sqrt{Nd})$ for $\lambda = O(\sqrt{Nd})$ [2067, 2068]. A logarithmic depth is only reachable with $O(N)$ parallel ancillas, so depth and ancilla count cannot both be logarithmic.

It is a quantum circuit created for a specific dataset that is injected in specific parts of a quantum algorithm circuit, and usually in a repetitive manner, for the many unitary exponentiations required by a QPE. Each query costs $O(N)$ T gates, a Toffoli costing 7 T gates in the provably optimal ancilla-free Clifford+T decomposition, or 4 T gates using the measurement-based temporary AND construction whose uncomputation is Clifford-only [2069, 2070].

Algorithm papers routinely write “assume QRAM access in $O(\log N)$ time” and charge $O(\log N)$ to the complexity. End-to-end accounting with a QROM implementation instead charges $O(N)$ T gates per access. This discrepancy is often what separates the claimed exponential speedup from the actual resource cost after correct accounting. But QROM can be optimized in multiple ways, particularly to reduce the number of Toffoli and T gates [2071, 2072].

HHL as published in 2009 requires $O(\kappa^2/\epsilon)$ applications of the matrix oracle, where κ is the condition number and ϵ the target precision [187]. That scaling has since been improved: variable-time amplitude amplification makes the κ dependence near-linear, and the Chebyshev and Fourier decompositions of Childs, Kothari and Somma replace the $1/\epsilon$ factor by $\text{polylog}(1/\epsilon)$ [2073].

Any honest cost accounting must use the improved figure, not the 2009 one.

For a representative example, take $n = 10^6$, $s = 6$ (a sparse FEM matrix, so $N = n \cdot s = 6 \times 10^6$ stored entries), $\kappa = 10^3$, $\varepsilon = 10^{-2}$ and $d = 32$ bits. With unary iteration the oracle costs $4N \approx 2.4 \times 10^7$ T gates, and with QROAM it drops to $O(\sqrt{Nd}) \approx 1.4 \times 10^4$ Toffoli, at the price of some 10^4 ancilla qubits. With near-optimal κ scaling the number of oracle calls is of order 10^4 to 10^5 rather than 10^8 . This gives:

$$2.4 \times 10^7 \times 5 \times 10^4 \approx 10^{12} \text{ T gates.}$$

With an optimistic FTQC rate of 10^6 T gates per second, this is around 10^6 seconds, roughly twelve days, landing to a few hours with QROAM. That is four orders of magnitude better than the naive accounting suggests, and it is still hopeless. The comparison point is not dense Gaussian elimination in $O(n^3)$, which nobody would run on a 6-sparse matrix, but conjugate gradient in $O(s \cdot n \cdot \sqrt{\kappa}) \approx 2 \times 10^8$ operations, a fraction of a second on a laptop. The gap is six to ten orders of magnitude in favor of the classical method, before even accounting for the fact that HHL returns a quantum state that cannot be read out in full [2057]. The exponential speedup in n is real asymptotically, and under the assumptions used in this example, data access and readout costs remove the practical advantage at the problem sizes considered here. How far that conclusion generalizes depends on the input access model, the conditioning, the sparsity, the required precision, the output observable and the classical baseline.

QROM can also be associated with the process of Quantum Mass Production to efficiently use a stored unitary and exponentiate it several times, as in a quantum phase estimation algorithm [2074, 2075].

Table 11 summarizes the key distinctions between qRAM and QROM.

Other quantum memories

- **RQAM (Random Quantum Access Memory)** or Random Access Quantum Memory (RAQM) which can encode arbitrary data in single of multiple qubits [2076–2078].
- **Probabilistic Quantum Memory (PQM)** stores and simultaneously analyzes r patterns while using only n qubits. A quantum computer therefore would need $O(n)$ qubits as opposed to $O(rn)$ bits of associative memory on a classical computer [2079, 2080].
- **Quantum Associative Memory (QAM)** employs the dynamics of open quantum systems to store and retrieve associative patterns, where the memory responds to an input by recalling a related stored pattern, even when the input is corrupted or incomplete due to noise [2033, 2081].

- **QKD repeater memory.** Some quantum memory is also required in quantum key distribution repeaters [2082, 2083]. We cover this in the communication section from this book.
- **Quantum sensing memory** can be useful in various situations like with quantum sensing and for creating deterministic sources of photons [2084].

Quantum memory has the same shortcomings as computing qubits: it is subject to noise, decoherence and correlated errors [2085]. As a result, quantum error correction codes are necessary to protect these memories over time, particularly as they are ingesting significant amounts of data [2056, 2086–2088]. So, we will also see logical qubits being built out of physical qubits in the various varieties of quantum memories [2089]. The QEC codes capabilities depend on the function of the quantum memory. If it ingests data that is prepared elsewhere, QEC will not need to support T gates and its related costly magic state distillation [2090, 2091]. If arbitrary states must be synthesized inside the memory, non-Clifford resources become necessary. There are also quantum error mitigation techniques (QEM) available to correct bucket-brigade qRAM [2092].

On the architecture side, many innovations are in the making to support efficient qRAM [2094–2096], like the 2025 theoretical blueprint from Yale to pipeline multiple quantum queries simultaneously while maintaining desirable scalings in query speed and fidelity. Their superconducting-qubit “Fat-Tree” QRAM would perform $O(\log N)$ independent queries in $O(\log N)$ time using $O(N)$ qubits, bringing parallelism compared with regular QRAM architectures [2093]. Query parallelization in a bucket brigade had been investigated earlier by Alexandru Paler et al. [2097]. The architecture uses “routers” which interpret part of the memory address and direct the query to the correct branch of the bucket brigade tree, all this being implemented with regular quantum gates (Figure 291).

5.3.4 Quantum memory physical implementations

No demonstrated quantum memory architecture yet meets the scale, fidelity, lifetime, bandwidth and FTQC requirements of a general-purpose computing qRAM. Experimental memories do exist and some are described below, including small bucket-brigade prototypes. Research is making progress, with targeted use cases that are more related to secure telecommunications and for quantum optical repeaters. At this stage, the advent of qRAM for quantum computing is more difficult to predict than scalable quantum computing! The most promising quantum memory technologies are coupling cold atoms and photon polarization [2098]. I mention here only the quantum memory adapted to quantum computing.

Approach	Query time	Hardware qubits	T-gates/query	Error correction
Hardware QRAM	$O(\log N)$	$O(N)$ coherent	none (native routing)	Needed at high query counts
QROM, unary iteration	$O(N)$ depth	$O(\log N)$ ancilla	$4N - 4$	Standard FTQC
QROM, QROAM	$O(N/\lambda)$ depth	$O(\lambda d)$ ancilla	$O(\sqrt{Nd})$	Standard FTQC

Table 11: QRAM vs. QROM resource comparison for N data items. Each row describes one named construction under its own cost conventions and not a general property of the approach. The T counts are logical counts and exclude Clifford gates, and the hardware QRAM row having no logical T gate does not mean it is free: native routing still requires physical operations, coherent idling of the router qubits and, above a certain query count, its own error correction overhead.




Figure 291: The parallelized Fat-Tree bucket brigade qRAM pipeline proposed in the 2025 Yale blueprint [2093].

TABLE OF COMPARISON BETWEEN DIFFERENT QRAM TECHNOLOGIES.

Feature	Bucket-Brigade QRAM [36]	Fanout QRAM [37]	Flip-Flop QRAM [44]	Qudits-based memory [48]	Approximate PQC-based [49] & EQGAN QRAM [46]
Structure	Bifurcation graph	Bifurcation graph	Quantum circuit	Higher states	Parametric Quantum Circuit
Circuit width ($n = \# \text{address lines}$)	$O(2^n)$	$O(2^n)$	$O(n)$	Dependent on d (# qudit states)	$O(n)$
Circuit depth ($n = \# \text{address lines}$)	$O(2^n)$	$O(2^n)$	$O(2^n)$	Dependent on d (# qudit state)	$O(1)$
Unique qualities	Qubits are routed in a sequential fashion	Qubits controlling exponential quantum switches	Quantum circuit-based	Reduces requirements of ancilla qubits to 0	Trainable like a machine learning model
Implementation technology	Photons, trapped atoms	Photons, microwave cavities	Superconducting qubits, trapped ion qubits	Superconducting qudits, trapped ion qudits, OAM photonic qudits	Superconducting qubits, trapped ion qubits
Drawbacks	Exponential circuit width and depth	Exponential circuit width and depth, susceptible to decoherence	Exponential circuit depth	Unstable higher states	Performance degradation under noise (approx. QRAM), store only simple dataset (EQGAN)

Figure 292: Another segmentation of various forms of qRAM. Source: [2032]. Added in 2023.

In the part of this book on quantum communications, other physical implementations of quantum memory are mentioned that are suitable only for communications and the creation of entanglement resources between two systems.

- **Optical memories** are tested with ytterbium, a rare earth that can be controlled at high frequency [2099]. [2100] uses ytterbium nuclear spin in yttrium orthovanadate crystal (YVO_4 , V for vanadium) arranged in nanophotonic cavity. It stores polarization information in spin ensembles. Bell states are created with ytterbium and vanadium. Control is made with

675 and 991 MHz microwaves and optical readout at 984 nm. It operated at 460 mK. Other used materials are europium [2101] and YAG crystals [2102]. The process preserves the polarization of a single photon in a magnetic trap, rather for optical repeater applications in long-distance secure communication lines [2103].

- **Non-volatile switching** in an iron-germanium-tellurium van der Waals ferromagnet (Fe_3GeTe_2), toggled by local laser heating and operating near room temperature [2104]. This is a classical non-volatile memory effect and not a quantum memory as such, but it belongs to the broader set of candidate cryo-compatible memory technologies.
- **Photons** trapped in fiber cavities, the memory trapping photons in a low-loss cavity [2105].
- **GKP photonic states** in continuous variables interacting with atoms in cavities [2106, 2107].
- **NV centers** [2108, 2109] and other crystal defects are also tested [2110, 2111].
- **Atom donor spin** setups are another option for the storage of quantum states, like with bismuth in silicon [2112, 2113]. [2114] deals with using ensembles of bismuth donor spin in natural silicon, coupled to a planar superconducting niobium resonator, all operating at 100 mK with a resonant frequency of 7.093 GHz. Pulses are made of 1,200 photons. It seems to be used with individual qubits memory, not entangled qubits and amplitude encoding.
- **Trapped ions** as experimented in 2022 in China with 218 ions in a cryogenic Paul trap forming long chains and quasi-1D zigzag crystals, with storage coherence times of the order of hundreds of milliseconds on arbitrarily chosen edge ions, sympathetic cooling of a fraction of the ions in the middle of the chain being necessary to sustain it [2115].
- **Cold atoms ensembles** which store information from incoming photons and their property like polarization or orbital angular momentum. Light-matter interactions is usually based on Raman scattering, where a photon interacts with an atom ensemble and its collective excitation, and is absorbed, transferring its energy to the atom ensemble's internal state [2116].

Chinese scientists demonstrated in 2019 the storage of single-photon polarization qubits in laser-cooled rubidium ensembles held in magneto-optical traps, using balanced two-channel electromagnetically induced transparency, with a storage and retrieval efficiency above 85% and a fidelity above 99%, the single-channel efficiency reaching 90.6% [2117] (Figure 293). Rubidium atoms were cooled with lasers to 200 μK .

The same year, another team in China reported a 105-qubit memory using 210 memory cells and a dual-rail representation of photonic qubits, with fidelities around 90% [2118]. What the paper established is a number of modes, an encoding, storage time, efficiency and fidelity. Entanglement across the stored qubits was not demonstrated, which is not the same as showing it cannot be reached, and storing a 2^N -amplitude state vector is not an operational benchmark for such a device. Other techniques are based on cesium with fidelities reaching 99% [2098, 2119].

This is also the technique developed by Julien Laurat at ENS LKB in Paris and implemented by Weling for both quantum repeaters and quantum computers interconnect. A related work in Canada is dynamically controlling rubidium's transparency to trap single photons [2120]. In practice, photons are stored for a thousandth of a second, but this would be sufficient for optical telecommunication repeaters. Another work in France from the Pasqal team used cold atoms to store quantum information [2121].

Variants exist with atomic ensemble memories using an Atomic Frequency Comb (AFC) protocol implemented in either a two- or three level system, referred to as optical and spin wave storage, respectively [2122]. And a 25-dimensional orbital angular momentum mode photon storage in cold atom ensembles in China, with a storage efficiency close to 60% and a fidelity of 84.2%, uniform across the 25 spatial modes. That efficiency is modest in absolute terms but is the highest reported so far for a high-dimensional quantum memory, earlier attempts being far below [2123].

In a hybrid approach, the University of Cambridge stores some quantum bit information in an electron spin hidden in haystack of 100,000 atom nuclei. The electron spin and the whole haystack are controlled by a laser. But the nuclei surrounding the electron make it difficult to entangle several qubits [2124]!

- **Nuclear spins** qubits storage [2125] like with ionized $^{31}\text{P}^+$ donors in isotopically purified ^{28}Si , where a 39-minute storage time was reached at room temperature back in 2013 [2126].
- **Superconducting cavities**, as explored at Yale [2127] and in Russia [2128].
- **Transmon-based bucket brigade routers**. Coherent quantum routers, the four-body building block of a bucket brigade tree, were demonstrated on superconducting processors in 2025 by the Guo group in China [2129] and at Chicago [2130]. In the same year, a Zhejiang University team ran the first end-to-end bucket-brigade qRAM on a programmable superconducting processor, with the full cycle of address loading, data loading, data writing, data retrieval and address retrieval [2055]. Query fidelities reached 0.800 ± 0.026 for two layers and 0.604 ± 0.005 for




Figure 293: A cold atom ensemble based single qubit memory, using rubidium atoms laser-cooled to $200 \mu\text{K}$ in a magneto-optical trap. Source: [2117].





Figure 294: A circuit implementing a 3-layer bucket-brigade qRAM with superconducting qubits. Source: [2055].

three layers, both after post-selection based error mitigation on the router qubits, against 0.694 raw for the two-layer case. The three-layer device used 16 qubits. This is a hardware experiment, not a numerical simulation, but the routers are implemented digitally, decomposed into native single-qubit and CZ gates rather than realized as a native four-body interaction, with quantum teleportation used to map the binary tree onto a 2D square grid without losing the logarithmic time complexity.

- **Quantum mechanical oscillators** coupled with superconducting qubits, although not necessarily to build some qRAM [2131].

- **Bosonic superconducting qubits** can implement passively corrected quantum memory [2132, 2133].
- **Dual-rail superconducting qubits and surface acoustic wave (SAW) phonons** as studied at the University of Chicago et Argonne National Laboratory [2134].
- **Quantum memcapacitors**, studied in China and Germany [2135].
- Memory with **error correction** using the honeycomb technique or Floquet codes [2136].



5.3.5 Bibliography and notes

These are the bibliographical references and notes for the **Quantum memory** subsection.

- [187] [Quantum algorithm for linear systems of equations](#), by Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd, *Physical Review Letters*, October 2009 (4 pages).
- [2028] [Reducing Refresh Overhead with In-DRAM Error Correction Codes](#), by Hanbyeol Kwon et al., *International SoC Design Conference, ISOC*, October 2021 (4 pages) 📄.
- [2029] [Quantum random access memory](#), by Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone, *arXiv*, March 2008 (4 pages).
- [2030] [Quantum Memory: A Missing Piece in Quantum Computing Units](#), by Chenxu Liu, Ang Li, et al., *arXiv*, November 2023 (41 pages).
- [2031] [Circuit-Based Quantum Random Access Memory for Classical Data](#), by Daniel K. Park, Francesco Petruccione, and June-Koo Kevin Rhee, *Scientific Reports*, March 2019 (8 pages).
- [2032] [Quantum Random Access Memory For Dummies](#), by Koushtubh Phalak, Avimita Chatterjee, and Swaroop Ghosh, *arXiv*, May 2023 (12 pages).
- [2033] [Theoretical framework for quantum associative memories](#), by Adrià Labay-Mora, Eliana Fiorelli, Roberta Zambrini, and Gian Luca Giorgi, *arXiv*, August 2024 (2 pages).
- [2034] [From Bits to Qubits: The Theory and Practice of Quantum Data Encoding](#), by Xiao-Ming Zhang, Bálint Koczor, Xiao Yuan, et al., *arXiv*, September 2026 (72 pages).
- [2035] [Quantum Machine Learning and qRAM](#), by Behnam Kia, 2018 (59 slides).
- [2036] [Quantum Algorithms for Linear Algebra and Machine Learning](#), by Anupam Prakash, 2014 (91 pages).
- [2037] [Quantum 101: Do I need a quantum RAM?](#), by Olivia Di Matteo, May 2020 (58 slides).
- [2038] [Architectures for a quantum random access memory](#), by Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone, *arXiv*, November 2008 (10 pages).
- [2039] [QRAM: A Survey and Critique](#), by Samuel Jaques and Arthur G. Rattew, *arXiv*, May 2023 (38 pages).
- [2040] [Quantum algorithm for finding minimum values in a Quantum Random Access Memory](#), by Anton S. Albino, Lucas Q. Galvão, Ethan Hansen, Mauro Q. Nooblat Neto, and Clebson Cruz, *arXiv*, January 2023 (6 pages).
- [2041] [Novel oracle constructions for quantum random access memory](#), by Ákos Nagy and Cindy Zhang, *arXiv*, June 2024 (18 pages).
- [2042] [A resource-efficient quantum-walker Quantum RAM](#), by Giuseppe De Riso, Giuseppe Catalano, Seth Lloyd, Vittorio Giovannetti, and Dario De Santis, *arXiv*, August 2025 (15 pages).
- [2043] [Efficient Quantum State Preparation with Bucket Brigade QRAM](#), by Alessandro Berti and Francesco Ghisoni, *arXiv*, October 2025 (29 pages).
- [2044] [Exponential separations between learning with and without quantum memory](#), by Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li, *arXiv*, November 2021 (77 pages).
- [2045] [On the robustness of bucket brigade quantum RAM](#), by Srinivasan Arunachalam, Vlad Gheorghiu, Tomas Jochym-O'Connor, Michele Mosca, and Priyaa Varshinee Srinivasan, *New Journal of Physics*, December 2015 (20 pages).
- [2046] [Quantum Random Access Memory](#), by Aaron Green et al., 2019 (12 pages).
- [2047] [Methods for parallel quantum circuit synthesis, fault-tolerant quantum RAM, and quantum state tomography](#), by Olivia Di Matteo, 2019 (111 pages).
- [2048] [Fault Tolerant Resource Estimation of Quantum Random-Access Memories](#), by Olivia Di Matteo, Vlad Gheorghiu, and Michele Mosca, *IEEE Transactions on Quantum Engineering*, 2020 (13 pages).
- [2049] [Quantum memory based on concatenating surface codes and quantum Hamming codes](#), by Menglong Fang and Daiqin Su, *arXiv*, April 2025 (12 pages).
- [2050] [On the energy barrier of hypergraph product codes](#), by Guangqi Zhao, Andrew C. Doherty, and Isaac H. Kim, *arXiv*, July 2024 (8 pages).
- [2051] [Fundamental causal bounds of quantum random access memories](#), by Yunfei Wang, Yuri Alexeev, Liang Jiang, Frederic T. Chong, and Junyu Liu, *npj Quantum Information*, July 2024 (12 pages).
- [2052] [Efficient Complex-Valued State Preparation on Bucket Brigade QRAM](#), by Alessandro Berti and Francesco Ghisoni, *arXiv*, April 2026 (12 pages).
- [2053] [Quantum random access memory via quantum walk](#), by Ryo Asaka, Kazumitsu Sakai, and Ryoko Yahagi, *arXiv*, May 2021 (13 pages).
- [2054] [Resilience of Quantum Random Access Memory to Generic Noise](#), by Connor T. Hann, Gideon Lee, S. M. Girvin, and Liang Jiang, *PRX Quantum*, April 2021 (21 pages).
- [2055] [A bucket-brigade quantum random access memory](#), by Fanhao Shen, Yujie Ji, Debin Xiang, Yanzhe Wang, Ke Wang, Chuanyu Zhang, Chao Song, Liqiang Lu, Zhen Wang, and Jianwei Yin, *Nature Physics*, 2026 (8 pages) 📄 (preprint on *arXiv*).
- [2056] [A distillation-teleportation protocol for fault-tolerant QRAM](#), by Alexander M. Dalzell, András Gilyén, Connor T. Hann, Sam McArdle, Grant Salton, Quynh T. Nguyen, Aleksander Kubica, and Fernando G.S.L. Brandão, *arXiv*, May 2025 (49 pages).
- [2057] [Read the fine print](#), by Scott Aaronson, *Nature Physics*, April 2015 (3 pages) 📄.
- [2058] [Quantum simulation of chemistry with sublinear scaling in basis size](#), by Ryan Babbush, Dominic W. Berry, Jarrod R. McClean, and Hartmut Neven, *npj Quantum Information*, 2019.
- [2059] [Simulating Sparse Hamiltonians with Star Decompositions](#), by Andrew M. Childs, Robin Kothari, and Rolando D. Somma, *Theory of Quantum Computation, Communication, and Cryptography (TQC 2010)*, 2010 📄.
- [2060] [Operational Framework for a Quantum Database](#), by Carla Rieger, Michele Grossi, Gian Giacomo Guerreschi, Sofia Vallecorsa, and Martin Werner, *arXiv*, February 2025 (27 pages).
- [2061] [Efficient Classical Simulation of the Quantum Random Access Memory](#), by Yun-Jie Wang, Guo-Ping Guo, et al., *arXiv*, March 2025 (11 pages).
- [2062] [Resource-efficient simulation of noisy quantum circuits and application to network-enabled QRAM optimization](#), by Luis Bugalho, Emmanuel Zambrini Cruzeiro, Kevin C. Chen, Wenhan Dai, Dirk Englund, and Yasser Omar, *npj Quantum Information*, October 2023 (11 pages).
- [2063] [Optimal \(controlled\) quantum state preparation and improved unitary synthesis by quantum circuits with any number of ancillary qubits](#), by Pei Yuan and Shengyu Zhang, *arXiv*, May 2023 (21 pages).
- [2064] [Data centers with quantum random access memory and quantum networks](#), by Junyu Liu, Connor T. Hann, and Liang Jiang, *arXiv*, September 2023 (23 pages).

- [2065] [Circuit-based quantum random access memory for classical data with continuous amplitudes](#), by Tiago M. L. de Veras, Ismael C. S. de Araujo, Daniel K. Park, and Adenilton J. da Silva, arXiv, November 2020 (11 pages).
- [2066] [Optimization of Quantum Read-Only Memory Circuits](#), by Koustubh Phalak, Mahabubul Alam, Abdullah Ash-Saki, Rasit Onur Topaloglu, and Swaroop Ghosh, arXiv, April 2022 (6 pages).
- [2067] [Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity](#), by Ryan Babbush, Craig Gidney, Dominic W. Berry, Nathan Wiebe, Jarrod McClean, Alexandru Paler, Austin Fowler, and Hartmut Neven, Physical Review X, October 2018 (36 pages)  (preprint on [arXiv](#)).
- [2068] [Trading T gates for dirty qubits in state preparation and unitary synthesis](#), by Guang Hao Low, Vadym Kliuchnikov, and Luke Schaeffer, Quantum, June 2024 (26 pages).
- [2069] [An algorithm for the T-count](#), by David Gosset, Vadym Kliuchnikov, Michele Mosca, and Vincent Russo, Quantum Information and Computation, 2014 (20 pages).
- [2070] [Halving the cost of quantum addition](#), by Craig Gidney, Quantum, June 2018 (6 pages).
- [2071] [Halving the cost of QROM](#), by Danial Motlagh and Matthew Pocrnic, arXiv, 2026 (8 pages).
- [2072] [Optimal T Counts under Sparsity: from QROM to State Preparation and Block Encoding](#), by Tongyang Li, Fengning Ou, Xinzhaoh Wang, Penghui Yao, Pei Yuan, and Shengyu Zhang, arXiv, July 2026 (46 pages).
- [2073] [Quantum Algorithm for Systems of Linear Equations with Exponentially Improved Dependence on Precision](#), by Andrew M. Childs, Robin Kothari, and Rolando D. Somma, SIAM Journal on Computing, 2017 (31 pages)  (preprint on [arXiv](#)).
- [2074] [Quantum Mass Production Theorems](#), by William Kretschmer, arXiv, May 2023 (12 pages).
- [2075] [Productionizing Quantum Mass Production](#), by William J. Huggins, Tanuj Khattar, and Nathan Wiebe, arXiv, May 2025 (34 pages).
- [2076] [Realization of a programmable multi-purpose photonic quantum memory with over-thousand qubit manipulations](#), by Sheng Zhang, Jixuan Shi, Zhaibin Cui, Ye Wang, Yukai Wu, Luming Duan, and Yunfei Pu, arXiv, April 2024 (17 pages).
- [2077] [A Cascaded Random Access Quantum Memory](#), by Ziqian Li, Eesh Gupta, Fang Zhao, Riju Banerjee, Yao Lu, Tanay Roy, Andrew Oriani, Andrei Vrajitoarea, Srivatsan Chakram, and David I. Schuster, arXiv, March 2025 (24 pages).
- [2078] [Quantum Storage of Qubits in an Array of Independently Controllable Solid-State Quantum Memories](#), by Markus Teller, Susana Plascencia, Samuele Grandi, and Hugues de Riedmatten, arXiv, September 2025 (11 pages).
- [2079] [Probabilistic Quantum Memories](#), by Carlo A. Trugenberger, arXiv, July 2001 (4 pages).
- [2080] [EP-PQM: Efficient Parametric Probabilistic Quantum Memory with Fewer Qubits and Gates](#), by Mushahid Khan, Jean Paul Latyr Faye, Udson C. Mendes, and Andriy Miranskyy, arXiv, April 2022 (27 pages).
- [2081] [High-capacity associative memory in a quantum-optical spin glass](#), by Brendan P. Marsh, Benjamin L. Lev, et al., arXiv, September 2025 (6 pages).
- [2082] [One-hour coherent optical storage in an atomic frequency comb memory](#), by Yu Ma, You-Zhi Ma, Zong-Quan Zhou, Chuan-Feng Li, and Guang-Can Guo, Nature Communications, April 2021.
- [2083] [Space-borne quantum memories for global quantum communication](#), by Mustafa Gündoğan, Markus Krutzik, et al., arXiv, June 2020 (11 pages).
- [2084] [Quantum Memories. A Review based on the European Integrated Project “Qubit Applications \(QAP\)”](#), by C. Simon, Harald Weinfurter, J. Wrachtrup, R.J. Young, et al., arXiv, March 2010 (22 pages).
- [2085] [Effect of Correlated Errors on Quantum Memory](#), by Smita Bagewadi and Avhishek Chatterjee, arXiv, February 2025 (7 pages).
- [2086] [Heterogeneously error-corrected QRAMs](#), by Ansh Singal and Kaitlin N. Smith, arXiv, April 2025 (10 pages).
- [2087] [Analysis and Suppression of Errors in Quantum Random Access Memory under Extended Noise Models](#), by Rohan Mehta, Gideon Lee, and Liang Jiang, arXiv, December 2024 (17 pages).
- [2088] [Fast and Error-Correctable Quantum RAM](#), by Francesco Cesa, Hannes Bernien, and Hannes Pichler, arXiv, March 2025 (40 pages).
- [2089] [Maximizing the Yield of Bucket Brigade Quantum Random Access Memory using Redundancy Repair](#), by Dongmin Kim, Sovanmonynuth Heng, Sengthai Heng, and Youngsun Han, arXiv, May 2024 (12 pages).
- [2090] [Fault-Tolerant Quantum Memory using Low-Depth Random Circuit Codes](#), by Jon Nelson, Gregory Bentsen, Steven T. Flammia, and Michael J. Gullans, arXiv, November 2023 (18 pages).
- [2091] [A quantum random access memory \(QRAM\) using a polynomial encoding of binary strings](#), by Priyanka Mukhopadhyay, Scientific Reports, March 2025.
- [2092] [Error-Mitigated Quantum Random Access Memory](#), by Wenbo Shi, Neel Kanth Kundu, Matthew R. McKay, and Robert Malaney, arXiv, October 2024 (8 pages).
- [2093] [Fat-Tree QRAM: A High-Bandwidth Shared Quantum Random Access Memory for Parallel Queries](#), by Shifan Xu, Alvin Lu, and Yongshan Ding, ASPLOS 2025, Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2025 (17 pages)  (preprint on [arXiv](#)).
- [2094] [Systems Architecture for Quantum Random Access Memory](#), by Shifan Xu, Connor T. Hann, Ben Foxman, Steven M. Girvin, and Yongshan Ding, MICRO 2023, Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture, 2023 (13 pages)  (preprint on [arXiv](#)).
- [2095] [Approximate Quantum Random Access Memory Architectures](#), by Koustubh Phalak, Junde Li, and Swaroop Ghosh, arXiv, October 2022 (5 pages).
- [2096] [Scalable and High-Fidelity Quantum Random Access Memory in Spin-Photon Networks](#), by Kevin C. Chen, Wenhan Dai, Carlos Errando-Herranz, Seth Lloyd, and Dirk Englund, PRX Quantum, July 2021 (17 pages).
- [2097] [Parallelizing the queries in a bucket-brigade quantum random access memory](#), by Alexandru Paler, Oumarou Oumarou, and Robert Basmadjian, Physical Review A, September 2020 (9 pages)  (preprint on [arXiv](#)).
- [2098] [Highly-efficient quantum memory for polarization qubits in a spatially-multiplexed cold atomic ensemble](#), by Pierre Vernaz-Gris, Kun Huang, Mingtao Cao, Alexandra S. Sheremet, and Julien Laurat, Nature Communications, January 2018 (7 pages).
- [2099] [Multimode and Random-Access Optical Quantum Memory via Adiabatic Phase Imprinting](#), by Nasser Gohari Kamel, Daniel Oblak, et al., arXiv, June 2025 (12 pages).
- [2100] [Nuclear spin-wave quantum register for a solid-state qubit](#), by Andrei Ruskuc, Nature, February 2022 (32 pages)  (preprint on [arXiv](#)).
- [2101] [On-Demand Quantum Storage of Photonic Qubits in an On-Chip Waveguide](#), by Chao Liu, Physical Review Letters, December 2020 (6 pages)  (preprint on [arXiv](#)).
- [2102] [Machine-Learning-Enhanced Quantum Optical Storage in Solids](#), by Yisheng Lei, Haechan An, Zongfeng Li, and Mahdi Hosseini, arXiv, April 2024 (5 pages).
- [2103] [Simultaneous coherence enhancement of optical and microwave transitions in solid-state electronic spins](#), by An-

- tonio Ortu et al., *Nature Materials*, December 2017 (10 pages)  (preprint on [arXiv](#)).
- [2104] [Reversible non-volatile electronic switching in a near-room-temperature van der Waals ferromagnet](#), by Han Wu, Ming Yi, et al., *Nature Communications*, March 2024.
- [2105] [Toward a Quantum Memory in a Fiber Cavity Controlled by Intracavity Frequency Translation](#), by Philip J. Bustard, Kent Bonsma-Fisher, Cyril Hnatovsky, Dan Grobnc, Stephen J. Mihailov, Duncan England, and Benjamin J. Sussman, *arXiv*, March 2022 (7 pages).
- [2106] [Entangling Quantum Memories at Channel Capacity](#), by Prajit Dhara, Liang Jiang, and Saikat Guha, *arXiv*, October 2024 (16 pages).
- [2107] [Interfacing Gottesman-Kitaev-Preskill Qubits to Quantum Memories](#), by Prajit Dhara, Liang Jiang, and Saikat Guha, *arXiv*, October 2024 (17 pages).
- [2108] [Storing quantum information in spins and high-sensitivity ESR](#), by John Morton et al., *Journal of Magnetic Resonance*, September 2017 (12 pages).
- [2109] [Electronic Noise Considerations for Designing Integrated Solid-State Quantum Memories](#), by Tzu-Yung Huang, David A. Hopper, Kaisarbek Omirzakhov, Mohamad Hossein Idjadi, S. Alexander Breitweiser, Firooz Aflatouni, and Lee C. Bassett, *arXiv*, December 2024 (13 pages).
- [2110] [A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute](#), by C. E. Bradley et al., 2019 (12 pages).
- [2111] [Multiplexed control of spin quantum memories in a photonic circuit](#), by D. Andrew Golter, Gerald Gilbert, Dirk Englund, et al., *arXiv*, April 2023 (10 pages).
- [2112] [Strong coupling of a superconducting flux qubit to single bismuth donors](#), by T. Chang, I. Holzman, S. Q. Lim, D. Holmes, B. C. Johnson, D. N. Jamieson, and M. Stern, *arXiv*, November 2024 (22 pages).
- [2113] [Addressing spins at the clock transitions with a frequency- and bandwidth-tunable superconducting resonator](#), by Yutian Wen, V. Ranjan, T. Lorriaux, D. Vion, B. Huard, A. Bienfait, E. Flurin, and P. Bertet, *arXiv*, October 2025 (11 pages).
- [2114] [Random-access quantum memory using chirped pulse phase encoding](#), by James O'Sullivan, Oscar W. Kennedy, Kamanasish Debnath, Joseph Alexander, Christoph W. Zollitsch, Mantas Šimėnas, Akel Hashim, Christopher N. Thomas, Stafford Withington, Irfan Siddiqi, Klaus Mølmer, and John J.L. Morton, *arXiv*, June 2022 (27 pages).
- [2115] [Experimental realization of a multiqubit quantum memory in a 218-ion chain](#), by R. Yao, W.-Q. Lian, Y.-K. Wu, G.-X. Wang, B.-W. Li, Q.-X. Mei, B.-X. Qi, L. Yao, Z.-C. Zhou, L. He, and L.-M. Duan, *Physical Review A*, December 2022 (7 pages).
- [2116] [Enhancing Quantum Memories with Light-Matter Interference](#), by Paul M. Burdekin, Ilse Maillette de Buy Wenniger, Stephen Sagona-Stophel, Jerzy Szuniewicz, Aonan Zhang, Sarah E. Thomas, and Ian A. Walmsley, *arXiv*, December 2024 (30 pages).
- [2117] [Efficient quantum memory for single photon polarization qubits](#), by Yunfei Wang, Jianfeng Li, Shanchao Zhang, Keyu Su, Yiru Zhou, Kaiyu Liao, Shengwang Du, Hui Yan, and Shi-Liang Zhu, *Nature Photonics*, May 2019 (6 pages)  (preprint on [arXiv](#)).
- [2118] [Experimental realization of 105-qubit random access quantum memory](#), by N. Jiang, Y.-F. Pu, W. Chang, C. Li, S. Zhang, and L.-M. Duan, *npj Quantum Information*, April 2019.
- [2119] [Efficient reversible entanglement transfer between light and quantum memories](#), by M. Cao, F. Hoffet, S. Qiu, A. S. Sheremet, and J. Laurat, *arXiv*, April 2021 (6 pages).
- [2120] [Coherent storage and manipulation of broadband photons via dynamically controlled Autler-Townes splitting](#), by Erhan Saglamyurek, Taras Hrushevskiy, Anindya Rastogi, Khabat Heshami, and Lindsay J. LeBlanc, *arXiv*, October 2017 (14 pages).
- [2121] [Storage and Release of Subradiant Excitations in a Dense Atomic Cloud](#), by Giovanni Ferioli et al., *PRX*, May 2021 (12 pages).
- [2122] [Lineshape Optimization in Inhomogeneous A-type Quantum Memory](#), by Kai Shinbrough, Donny R. Pearson Jr., Virginia O. Lorenz, and Elizabeth A. Goldschmidt, *arXiv*, May 2024 (8 pages).
- [2123] [Highly Efficient Storage of 25-Dimensional Photonic Qudit in a Cold-Atom-Based Quantum Memory](#), by Ming-Xin Dong, Guang-Can Guo, Dong-Sheng Ding, Bao-Sen Shi, et al., *Physical Review Letters*, December 2023 (7 pages)  (preprint on [arXiv](#)).
- [2124] [Quantum sensing of a coherent single spin excitation in a nuclear ensemble](#), by Daniel M. Jackson, Dorian A. Gangloff, Jonathan H. Bodey, Leon Zaporski, Clara Bachorz, Edmund Clarke, Maxime Hugues, Claire Le Gall, and Mete Atatüre, *arXiv*, August 2020 (21 pages).
- [2125] [Nuclear Spin Quantum Memory in Silicon Carbide](#), by Benedikt Tissot, Michael Trupke, Philipp Koller, Thomas Astner, and Guido Burkard, *arXiv*, August 2022 (12 pages).
- [2126] [Room-Temperature Quantum Bit Storage Exceeding 39 Minutes Using Ionized Donors in Silicon-28](#), by Kamyar Saeedi, Stephanie Simmons, Jeff Z. Salvail, Phillip Dluhy, Helge Riemann, Nikolai V. Abrosimov, Peter Becker, Hans-Joachim Pohl, John J. L. Morton, and Mike L. W. Thewalt, *Science*, November 2013 (4 pages)  (preprint on [arXiv](#)).
- [2127] [Quantum random access memory architectures using superconducting cavities](#), by D. K. Weiss, Shruti Puri, and S. M. Girvin, *arXiv*, March 2024 (21 pages).
- [2128] [Superconducting integrated on-demand quantum memory with microwave pulse preservation](#), by Aleksei R. Matanin, Ilya A. Rodionov, et al., *arXiv*, July 2025 (7 pages).
- [2129] [Demonstrating Coherent Quantum Routers for Bucket-Brigade Quantum Random Access Memory on a Superconducting Processor](#), by Sheng Zhang, Guo-Ping Guo, et al., *arXiv*, May 2025 (5 pages).
- [2130] [Implementation of a quantum addressable router using superconducting qubits](#), by Chuanhong Miao, Sebastien Leger, Ziqian Li, Gideon Lee, Liang Jiang, and David I. Schuster, *arXiv*, March 2025 (14 pages).
- [2131] [A mechanical quantum memory for microwave photons](#), by Alkim B. Bozkurt, Omid Golami, Yue Yu, Hao Tian, and Mohammad Mirhosseini, *arXiv*, December 2024 (30 pages).
- [2132] [Candidate for a passively protected quantum memory in two dimensions](#), by Simon Lieu, Yu-Jie Liu, and Alexey V. Gorshkov, *arXiv*, March 2023 (18 pages).
- [2133] [Superconducting Quantum Memory with a Suspended Coaxial Resonator](#), by Lev Krayzman, Chan U Lei, Suhas Ganjam, James Teoh, Luigi Frunzio, and Robert J. Schoelkopf, *arXiv*, February 2024 (11 pages).
- [2134] [Quantum random access memory with transmon-controlled phonon routing](#), by Zhaoyou Wang, Hong Qiao, Andrew N. Cleland, and Liang Jiang, *arXiv*, June 2025 (9 pages).
- [2135] [Microwave Quantum Memcapacitor Effect](#), by X.-Y. Qiu, F. Albarrán-Arriagada, et al., *arXiv*, May 2024 (6 pages).
- [2136] [A Fault-Tolerant Honeycomb Memory](#), by Craig Gidney, Michael Newman, Austin Fowler, and Michael Broughton, *arXiv*, December 2021 (17 pages).

This **Quantum memory** subsection contains 110 bibliographical references and notes containing at least 1,833 pages.

5.4 Energetics

The main motivation for creating quantum computers is their potential capacity to solve intractable problems in non-exponential times, or at least with significant acceleration in practical cases, even though not being exponential. This should make it possible to solve problems that are inaccessible to conventional HPC and the largest supercomputers. In some other cases, like with some NISQ architecture using quantum error mitigation, it will only be “just” somewhat faster or sometimes, provide better results, like with quantum machine learning or some many-body quantum physics simulations.

How this computing capacity translates in terms of energy consumption is a key open question, particularly as ICT (Information and Communication Technologies) energetic costs are skyrocketing, being currently amplified by the broad adoption of LLM-based consumer online tools.

This opens a Pandora’s box of questions that we’ll try to address in this part:

- **ICT lessons.** What is dissimilar and similar between classical computing and quantum computing when learning from History? What can we learn from the data? What went wrong, particularly with AI? Also, what is the real size of the AI power surge?
- **First assessment.** Have we already reached some energetic quantum advantage with the quantum supremacy experiments, starting with Google Sycamore in 2019? Contrarily, should we worry as some industry vendor roadmaps have started to build enormous systems, way larger than the first mainframes in the 1940s, or even, larger than the most powerful deployed exaflops supercomputers? Or should we believe other vendors like D-Wave, on the opposite side, who claim that their systems are energetically frugal? Do we have an idea of the potential energy consumption of large scale quantum computers, the one that will deliver on the promise to extend the reach of computing?
- **Cost vs value.** What can we do to estimate the size of the energetic problem or benefit of quantum computers? What will be their use cases and the related market size? The question is intertwined with the broad applications promises, the expectation to see the advent of some quantum “ChatGPT moment” and grandiose market forecasts. Responding to these questions deals with a mix of science and technology, applications and software, and economics. Sizing the market potential depends on the viable applications that will be created in the future and the volume of end-user organizations who will need it. Market volume estimates range between a subset of the supercomputing market and a large chunk of the cloud server segment!
- **Quantum computing energetics.** How is this discipline organized with its quantum level (with quantum thermodynamics and quantum energetics) and all macroscopic energetics, with enabling technologies? What are the scaling challenges with FTQC? How to assess the first energy resource estimates published, mostly in academic papers?
- **Energetic advantage.** What is the real source of a potential energetic advantage with quantum computing? Parallelization? Entanglement? Reversibility? Will quantum computing generate some potential energetic advantage for use cases that are still workable with classical solutions? Will use cases beyond that scale consume a reasonable amount of energy, particularly when they need thousands of logical qubits? We do not yet have the answer. It is a new research area worth the effort. How will we benchmark and optimize the energy consumption of quantum computers?
- **Power resource estimates.** How do you estimate the energy consumption of existing and future quantum computers and define a potential energetic advantage? That’s a tricky question for future systems which are only designed “on paper”. And it must also consider the significant cost of the classical systems that are necessary to run these QPUs. What is the current outcome from the academic world with regards to analyzing, estimating, and optimizing the energetic cost of quantum computers?
- **Qubit types differences.** Will there be some difference between the various qubit types and quantum computing paradigms like analog and gate-based quantum computers?
- **Actions.** What can be done to limit the pitfalls encountered with other technologies like AI? This is the realm of the *Quantum Energy Initiative* that I cofounded in 2022. It is still in its infancy. How scientists from the academic and industry vendor space are collaborating in this domain?

5.4.1 ICT lessons

The energetic footprint of all digital tools (information and communication technologies) is already significant. It was 5.3% of global electricity consumption in 2022, including manufacturing costs. The 6.1% for 2026 and 7% for 2030 are forecasts which combining various sources [2137, 2138] (Figure 295).

ICT energetic footprint tends to be focused on looking at the energetic consumption of products while they are used. It can be completed by carbon and greenhouse gas emission footprint analysis, which is embedding the whole lifecycle of ICT products [2145]. It leads to the concept of “carbon pricing” which encompasses a whole top-to-bottom lifecycle analysis of ICT products.



Figure 295: Global ICT energy footprint in 2022, 2026 and 2030, split between operational electricity (upward) and annualized manufacturing energy cost (downward), in TWh per year, with the equivalent continuous average power in GW on the right axis. The 2022 column is based on reported global data while the 2026 and 2030 columns are forecasts. The reconciliation uses similar methods like accounting for average and not peak power for power estimates. The data provides a lot of insights on the evolution of power consumption across multiple devices from end-user tools to cloud and data-center based servers. The sheer growth of LLM-based AI power consumption is staggering. Between 2022 and 2026, it went from 1.4 GW to 20 GW and is planned to reach 51 GW in 2030, provided there is no AI bubble burst in the interim. Sources for past and future global electricity consumption: [2137-2144]. © Olivier Ezratty, September 2024-September 2026

Since quantum computers are serverlike systems, we'll look at data centers energy consumption. Their functions can be categorized into at least five main segments:

Let's mention here the Green ICT working group launched in 2023 by Informatics Europe [2146]. They published a white paper on arXiv in April 2026 [2147]. Informatics Europe is an association that represents 200 university departments and public research institutes in computer science in Europe.

Regular Internet and cloud services

These use classical CPUs from Intel and AMD running very traditional workloads like web servers, email services, data storage, and other LAMP (Linux-Apache-MySQL-PHP) back-end services. That is where you have ecommerce sites, news sites, services from Google, Facebook, other social networks and online gaming. This category of computing will probably not be affected by quantum computing since none of these are computationally intractable problems.

AI workloads

This includes all LLM-based agents and chatbots, with their training and inference costs [2148]. Immense progress was made to optimize the energetic efficiency of LLM training. For example, Nvidia estimated in June 2024 that training GPT-4 with 1.8 tera parameters would have cost over 1 TWh with their 2016 Pascal GPU, 140 GWh with their 2017 Volta GPGPU, 40 GWh with their 2020 Ampere GPGPU and 3 GWh with their 2024 Blackwell GPGPU. The same keynote gave the matching inference figures, with the energy per generated token dropping by a factor 45,000 over eight years, from 17,000 J/token on Pascal to 0.4 J/token on Blackwell (Figure 296). Other multiple gains came from the usage of tensor networks to optimize LLM training and inferences [2149–2151].

One source estimated in 2023 that AI data center energy consumption could reach between 85 and 134 TWh in 2027 [2152]. We are already at about 175 TWh in 2026 (Figure 295)! Efficiency gains didn't indeed prevent the advent of Jevon's paradox and rebound effect as the competition is driving vendors to increase their systems complexity [2153, 2154] and model size [2155]. Also, LLM-generated code is expanding but the result is not necessarily energy efficient and new prompting methods must be invented to avoid this side effect [2156]. Inferences costs are growing as these LLM-based services have now several hundred million users [2157]. And LLM-based services are integrated everywhere from Google search services [2158, 2159] to Microsoft Office suite.

Thus the gigantic AI datacenters already deployed [2144] or announced. The 10 largest AI datacenters in the USA consumed 4.2 GW of power as of mid-2026 [2160]. Amazon planned to spend 16 GW by end-2027, as a cloud provider for third-party AI software and services com-

panies [2161]. It plans a 2.4 GW \$11B data center in Indiana [2162].

Meta plans a 5 GW datacenter and \$50B investment in Louisiana [2170]. xAI's Colossus 2 AI training data-center will need 2.2 GW in 2027 to run 550,000 Blackwell GPU fuelled by 59 natural gas turbines. It is the successor of the 300 MW Colossus 1.0. xAI has plans for 10 GW of AI compute power by the end of 2027. It is supposed to generate up to \$500B in revenue by end of 2027, which is a good sign of a financial AI bubble [2171, 2172]. OpenAI plans 9 GW of capacity by 2029 as part of the Stargate project announced in 2025. A 1.8 GW Joliet Technology Center is planned by Hillwood (real estate) and PowerHouse Data Centers (builder) in Illinois, representing a \$20B private investment, with long-term leasing to hyperscalers. In Europe, Mistral wants to build a 1 GW AI data-center in Europe by the end of 2030, with an interim 200 MW planned for 2027 [2173]. Nebius, a Dutch AI cloud company, wants to install 310 MW of computing power in Finland, on top of a 1 GW plan in Missouri [2174]. These few examples add about 20 GW of compute power to the whole cloud and colocation datacenters which was at about the same level in 2022 (Figure 295). China's data center capacity could reach 60 GW by 2030 [2175].




Figure 296: Nvidia estimates of the energy required to generate one inference token of a GPT-4 model, assuming a size of 1.8 tera parameters, a figure that was never officially disclosed by OpenAI, across its generations of GPGPU, from 17,000 J/token on Pascal down to 0.4 J/token on Blackwell, a 45,000× reduction in eight years. This model did not exist between 2016 and 2022, so the early points are retrospective estimates. The plot shows that more efficient GPUs enabled it. Nvidia also provided separate training figures for the same model, from over 1 TWh on the old Pascal GPU down to 3 GWh on the Blackwell GPU released in 2025. Source: [2176].

Despite the overpromises from vendors like IonQ [2177], FTQC will probably not be a solution, quantum computers being not appropriate for real time computing and large datasets loading.

Q The AI power consumption surge

The data exhibited in Figure 295 is based on multiple sources and some post-processing.

The 2022 column uses World Bank and ITU data [2137], with data-center including colocation, cloud and HPC, and telecom operators. Operational user devices and networks data come from the Malmudin et al. bottom-up model and its Ericsson update [2143, 2163]. DC 2026 and 2030 values follow the IEA base case [2142]. The AI share of 31% of DC power in 2026 comes from Gartner, and the 2030 AI share of 47% is consistent with the IEA statement that AI-focused facilities triple between 2025 and 2030 and with the JLL forecast that AI reaches half of installed DC capacity by 2030.

The 2022 AI figure is extrapolated from an estimate of 12 million AI GPUs consuming about 21 TWh at the end of 2023, then some 8% of data-center electricity [2163]. AI split use 2025 data from Google (43.6 TWh, +37%), Microsoft (37 TWh in FY2025, +24%) and Amazon (about 78 TWh for its data-center business), against about 72 TWh for the same group of companies in 2021.

HPC scientific computing is a small share of the data-center with power data coming from the TOP500 June 2026 list [2164] extrapolated to the listed systems that do not report power.

Manufacturing energy for servers, data-center and network hardware is estimated here using a top-down allocation using the 20 to 30% share of server lifecycle impact reported in product carbon footprints and in the LCA model of Wadenstein and Vanderbauwhede [2165], and a bottom-up count of roughly 14 million servers shipped per year including some 1.5 million AI servers [2165–2168]. As a comparable, semiconductor manufacturing electricity is projected at 237 TWh by 2030 [2169], and TSMC used 28.8 TWh in 2025. Device manufacturing data cover the full supply chain, including displays, batteries, boards and assembly, which is why they exceed the semiconductor industry electricity bill.

The datacenter GW data are average power and are lower than the GW announced capacities for vendors DCs, which usually run at about half of their capacity.

Surprisingly, in China, data-centers electricity demand surge is larger than the one coming from EV charging in China for the 2021-2027 period [2140].

Telecom networks equipment

These are located in their data centers and mobile base stations. Surprisingly, their energetic cost seemed higher than all other servers in the cloud, at least in 2022, before the LLM-surge.

The backbone of our landline (FTTH and DSL), on-demand television, and wireless communication has a huge cost. Only a few computing problems in that category are of interest for quantum computing, like network routing optimizations.

Bitcoin and cryptocurrency mining

The debate also raging about their potentially large energetic footprint, with various questionable comparison methodologies [2178, 2179]. For example, in the Blockchain realm, one may consider Ethereum's switch in September 2022 from proof-of-work to proof-of-stake for mining, with a significant energy saving of several orders of magnitude [2180]. Still, according to some estimates [2181], Bitcoin mining alone consumes about 100 TWh per year, a big chunk of the worldwide datacenter power consumption that was estimated to be 460 TWh in 2022 with expectations that it could grow to 1 PWh in 2026 [2182], with power (in W) being distributed on all continents, particularly in the USA, Europe and Asia

[2183] (it was 485 TWh in 2025 and projected to reach 950 TWh in 2030 by IEA in 2026 [2184]). There are some proposals to implement part of cryptocurrency mining with quantum computers to save energy, like BTQ Technologies (Taiwan, Canada) which wants to use a derivative of photonic boson sampling to implement energy-saving Bitcoin mining.

HPCs and supercomputers

These are usually dedicated to scientific computing tasks in aerospace engineering, weather forecasting and others, to machine and deep learning tasks for research, and will eventually be used to drive complex workloads involving quantum computers in so-called hybrid classical-quantum settings. The world's largest supercomputers have powers in the tens of MW (megawatts) range like Frontier from the DoE Oak Ridge National Laboratory in Tennessee and its 21 MW for 1.35 exaflops and 700 PB data storage, 9,408 AMD EPYC CPUs and 37,888 AMD Instinct MI250X GPUs. The most power-hungry system is Aurora, at the Argonne Leadership Computing Facility in Illinois, with 38.7 MW for 1.01 exaflops. New supercomputers are deployed every year, and we are entering the "exaflops" era.

The total mean power consumption of the Top 500 HPC worldwide can be estimated to be around 600 MW to 1 GW once the systems that do not report power are extrapolated, against about 420 MW for the documented systems alone [2164]. It would turn into 5 TWh of energy consumption per year, about 1% of the total

energy consumed by data centers worldwide and less than 0.4% of the global ICT energy consumption. It is quite marginal. This is the category of interest for quantum computing, which will supplement HPCs and supercomputers and, seemingly, work along them for many workloads. There are many ongoing efforts to optimize the energy efficiency of HPCs [2185, 2186]. This estimation is based on peak power for these supercomputers. The real consumption is probably about 60% to 70% of this level.

All these services use three main types of computing resources: “compute” resources for CPUs and GPUs, “data” resources for data storage and access using SSDs, hard disks and sometimes long-term tape storage, and “network” resources for connecting these systems to networks, the Internet and the cloud. All these hardware resources have an energetic footprint. In many situations, cooling data centers costs a lot, and many technological solutions are proposed to optimize this cost.

Servers and HPCs are installed “on premise” in the user’s own data centers, typically in defense, military and intelligence operations but also frequently in oil-and-gas, energy utilities and in some financial services institutions. Otherwise, these resources sit somewhere in the cloud either with hyperscalers like AWS and Microsoft or in rented datacenters like the ones from Equinix.

PUE *aka* Power Usage Effectiveness is a ratio dividing the total power consumed by a datacenter and the energy consumed by IT equipment (compute, storage, networking). The difference is mostly about the cost of cooling these equipment. The ratio is of about 1.05 to 1.1 for the best cases. It can even be optimized using machine learning techniques [2187]. The notion of PUE has not been defined yet for NISQ or FTQC systems given their specific cryogenic cooling requirements. And how would IT be defined here? It should include all classical computing and electronics/photronics involved in driving quantum processors.

Its distribution across countries and systems is detailed in Figure 297. The top consumer of MW is the USA with about 200 MW spread across 54 systems, the largest one consuming 38 MW at peak power. The distribution of power per system across the 200+ documented systems shows that we reach 1 MW at the 93rd position. There are only 9 supercomputers with a documented power consumption exceeding 10 MW. The average of these top 213 systems is 1.3 MW, for a documented total of about 280 MW. This power distribution could also be used to infer the potential QPU units market size based on their potential power consumption.

How will quantum computers fit in this picture? It depends on the stage of their maturity and commoditization. QPUs won’t replace regular Internet and cloud services and neither heavy AI-based workloads. We can expect that they will first sit in HPC centers. Since

QPU’s main benefit is providing some computing advantages compared to HPCs and supercomputers, they will probably become a subset of this market. Still, QPUs will complete HPCs more than replace them. Many of the existing HPC scientific applications are not suitable for quantum computing, like any digital simulation requiring large sets of data such as in weather forecasts or using the finite elements method and other methods to solve differential equations.




Figure 297: Chart showing how many MW are consumed by supercomputers in the world in June 2026, the values being in MW of peak power. The US is the top country for that respect with about 200 MW spread across 54 systems. And this is based on peak power consumption. The real power for the USA is probably around 120 to 150 MW. Data source: TOP500 [2164] and charting by Olivier Ezratty, September 2026.

When quantum computers scale up, they will mainly perform computations inaccessible to conventional supercomputers, like material and chemical simulations. But they will be fed by quantum circuits prepared classically by HPCs using traditional methods like DFT or various Hartree-Fock methods in the case of quantum chemical simulations, that are quite compute intensive. Since quantum computers will mostly sit in data centers and add another source of power drain to existing systems, it is interesting to look at classical comparables.

The energy efficiency of classical systems is the ratio of their performance to their energy consumption. Classical server efficiency can be expressed in FLOPS/W, where FLOPS is the number of floating-point operations per second. Since the birth of computing, this efficiency doubled about every 18 months, a trend known as Koomey’s law [2188]. Koomey himself later showed that the doubling time slowed to about 2.6 years after 2000, with the end of Dennard scaling, and an analysis of Green500 data over 2008-2023 lands at about 2.3 years [2189, 2190]. The DoE’s Frontier installed in Oak Ridge since 2022 in Tennessee has an energy efficiency of

52.93 GFLOPS/W with a computing power exceeding an exaflop. The more recent Venado system installed at the Los Alamos National Laboratory reaches an energy efficiency of 59.29 GFLOPS/W thanks to the use of recent Nvidia GPGPUs, the Grace Hopper GH200, but it is a much smaller system with “only” about 10 petaflops of Rmax computing power, for 15.6 petaflops of peak FP64 [2191].

QPU speed can be evaluated in CLOPS (circuit layer operations per second) which can’t be directly compared with FLOPS. CPU/QPU comparisons makes more sense with assessing their capability and resource to solve problems of given sizes with a given precision. It can use benchmarking tools like Eviden’s Q-Score for combinatorial problems or chemical simulations benchmarks. The closer quantum equivalent of GFLOPS/W could be a number of logical operations per second at a given logical fidelity per W (LOPS-f/W) but it is not yet an adopted figure of merit.

The supercomputing energetic efficiency gains have not prevented an explosion in global energy consumption to power digital technologies. Supercomputers account for a small share of the global digital energetic footprint, at about half a percent. It is increasing as usage grows. These phenomena are simply a new manifestation of the rebound effect, formalized by William Stanley Jevons in 1865. Efficiency gains automatically lead to a decrease in the cost of resources. Without regulation of markets and uses, they lead to an increase in global consumption. However, this does not mean that improving the energy efficiency of computers is inherently wrong. On the contrary, it is the only solution to maintain performance with limited energy and material resources. As we will see, the energetic cost of quantum computers will be highly variable across qubit types and vendors, and even among vendors in the same qubit type. Quantum computers will probably be installed in addition to existing classical systems rather than replacing them. Therefore, the rebound effect may not be observed per se.

One key under-tapped avenue to reduce classical computing energetic footprint lies with software optimizations and even the choice of programming languages [2192]. Unconventional classical computing options are also studied to save energy and covered in another section of this book, like reversible computing, superconducting computing and optical (non-quantum) computing [2193].

5.4.2 First assessments

At first glance, it looked like the energetic cost of quantum computing was several orders of magnitude lower than classical computers realizing the same task. It drove some interest for the energetics matter.

Quantum supremacy energetics

That was a (and my) naive interpretation of Google Sycamore’s 2019 quantum supremacy demonstration which did show a ratio of about one to one million in energy consumption compared to the IBM Summit supercomputer that was used as a comparison, and even when using the optimized algorithm and configuration proposed afterwards by IBM that ran in 2.5 days instead of the 10,000 years advertised by Google. But the benchmark was comparing apples and oranges with a randomized benchmark with no input data nor any useful output data. It was later shown by Xavier Waintal et al. that, when accounting for its high error rate and noise and using tensor networks, Sycamore’s performance could be emulated on a rather simpler classical server cluster [2194, 2195]. It reduced Google’s solution energy advantage to 1 to 30 ratio in favor of quantum computing as compared to a classical simulation. In 2023, Google itself did show that Sycamore 53-qubit randomized benchmarking could run in 6 seconds on the DoE Frontier classical supercomputer [236].

In 2024, a China research team led by Jian-Wei Pan classically simulated Google Sycamore’s original experiment 7 times faster and with 14 times less energy [2196, 2197]. The energetic advantage switched sides! Of course, this may be short lived. We will see if someday the same team reproduces the experiment simulating 72 or more qubits corresponding to Google’s Sycamore 2023 updated experiment. Some argue that as QPU power grows, any classical computing advantage will fade away, both in computing time and energy consumption, but not necessarily at the same moment.

In any case, this sort of comparison does not make a lot of sense since on both sides, we have some entirely useless computation with no input data. You cannot derive any sound conclusion for useful case studies. For example, data preparation or data loading with block-encoding, ansatz preparation or hybrid oracles may impose a significant toll on quantum computing performance for many use cases. Its avoidance in all supremacy experiments is a real shortcoming. Likewise, there are many other demonstrations of tensor network based contraction of NISQ circuit emulation overperforming QPUs, like the ones done around IBM’s quantum utility paper during Summer 2023 and afterwards. But they will probably reach some limits as QPUs grow in size and quality, potentially before the advent of FTQC QPUs.

On the other hand, another commonplace view is that the sheer power of about 15 kW that is required for cooling current small scale superconducting qubits processors is a showstopper. It gives the impression that quantum computers will be high-power consuming devices. This *may* not be true and forgets that a rack of Nvidia GPGPUs used for machine learning tasks has a power consumption above 30 kW. But how to compare these various beasts remains a new field of investigation.

NISQ energetics

With NISQ QPUs, the energy consumption of a quantum computer is relatively reasonable as shown in Table 12. A current quantum computer with superconducting qubits consumes about 25 kW, of which 16 kW comes from cryogenics. Cold atoms or photons quantum computers consume even less energy, because they do not require cryogenic cooling to 15 mK.

Photon qubits only require photon sources and detectors cooled at around 4K which is less energy hungry at face-value provided it scales well with the number of used qubits.

When thousands of qubits fit in these machines, their power consumption will increase due to the energetic cost of qubit control for initialization, quantum gates, error correction and qubit readout [2198]. Most qubits energetic costs come from the signals used for gate activations and readout.

These signals are microwave pulses, direct current pulses and laser beams. The related spent power seems to increase linearly with the number of qubits. But error correction requires a large number of physical qubits per logical qubit, adding another power consumption multiplying factor.

It will depend on the fidelity of the physical qubits and the ratio of physical qubits per logical qubit. The higher the fidelity, the lower this ratio will be. On top of that, the cryogenic cost grows very fast as the target temperature drops: the Carnot bound alone imposes a minimum input power of $(T_{\text{hot}}/T_{\text{cold}} - 1)$ watts per watt extracted, and real dilution refrigerators sit one to two orders of magnitude above it, so the wall-plug power needed to remove a given heat load roughly scales as $1/T^2$ at millikelvin temperatures. It also grows with the mass to cool down.

Creating a scalable quantum computer is clearly an optimization problem considering many energetic constraints, one being that the quantum computer's performance like the whole computing fidelity or accuracy reaches a given quantified target.

Qubits systems that operate at a cryogenic temperature are constrained by the cryostat cooling power and by the heat released within the cryostat. Superconducting and silicon spins qubits are the most challenging for that respect. Heat is generated by the inbound cable microwave attenuation filters and in the qubit readout related microwave amplifiers. In addition, the part of microwave generation and readout systems that is integrated into the cryostat have their own thermal footprint. All this must fit into the current thermal budget of the cryostats that is currently limited to 1W at the 4K stage and to about 25 μ W at the lower 15 mK stage.

The other way to be less constrained is to run the qubits at higher temperatures. This is what is possible with silicon spin qubits, which only require a temperature

between 100mK and 1K instead of 15 mK for superconductors. This increases the thermal budget for the control electronics at the qubit stage.

FTQC energetics

FTQC energetics are way different than with NISQ systems. Supporting large circuits in the megaqubit range and beyond requires quantum error correction to implement a logical memory that stays stable and coherent during the circuit execution, the realization of logical gates on this memory, and the generation and teleportation of magic states to target logical qubits. This has a cost that grows with the circuit size [2199]. When incurring all related costs, the physical overhead doesn't scale as fast as the circuit size but it is still significant.

The overhead of QEC/FTQC is twofold: physical with multiple orders of magnitude more physical qubits and all their enabling technologies, and time-wise since QEC bears its penalty here as well, due to the time-cost of repeated error syndrome measurements and classical error syndrome detection.

All this explains why the first power estimates for FTQC systems are much larger than with existing NISQ systems as shown in Table 12. The picture is however different across qubit modalities and even in some cases within qubit modalities. Photonic qubits exhibit the most expensive energetic cost with some vendors (PsiQuantum, Xanadu) while being reasonable for others like Quandela (not yet in the chart). Cold atoms, trapped ions and silicon qubits all sit in the low end of the power spectrum although no vendor has yet published any official data on this. Then, you must account for the speed and clock differences between these systems. If they scale well, silicon qubits will be much faster than cold atoms and trapped ions, and with a similar power cost.

Comparing classical and quantum computing

Real comparisons should be made with quantum computers solving real useful problems. It may potentially work with NISQ QPUs using quantum error mitigation following IBM's quantum utility mantra. It will more likely require FTQC architectures and a large number of physical qubits to implement error correction in a fault tolerant manner. Controlling these qubits uses energy-consuming conventional electronics and/or lasers. In that case, will quantum computers provide some energy advantage on top of a computing advantage, or do they risk turning into energy hogs [2216]? The same questions should be asked for other quantum technologies that could potentially be deployed at a large scale like quantum telecommunications and cryptography as well as quantum sensors.

In the industry vendors scene, the total power consumption of a quantum computer became a selling point, although not yet being perceived as being an important one. For example, AQT (trapped ions) explains that their 20-qubit system can be attached to a regular 220




Figure 298: Vendor projections of existing and future quantum processors power consumption. The FTQC points are vendor estimates based on their own assumptions about logical qubit definition, architecture, workload and facility boundary, while the classical supercomputer points are measured peak facility power available in the TOP500. The two families are not directly comparable. The chart's aim is to provide a comparison of the peak power consumed by existing HPC across a broad range of computing power as of June 2026 with existing NISQ and future FTQC systems. The related quantum computer workloads are indicated in logical qubits and circuit size. These estimates do not account for the cost of the supercomputers or HPCs that will drive these quantum computers for circuit preparation and any classical post-processing. The chart contains two indicative thresholds: one for 40 MW power, which corresponds to the largest current HPCs worldwide and considering the hypothesis that a maximum acceptable QPU power should be around that level. The second threshold is at 1 MW which corresponds to the 70th HPC in power consumption. This may lead to comparing QPUs using W (power) or \$ (hardware cost) per number of logical quantum operations per second (lqops), the smaller divider being the better, and for various levels of lqops (megaqops, gigaqops, teraqops, ...), given the FTQC overhead and thus power and hardware costs depend on the number of operations. This would account for the speed differences between QPUs. The chart provides some indication on how FTQC power consumption could potentially impact their market size when using HPC comparables in the grey rectangles at the top. Vendor data may change over time as it depends on pending optimizations, particularly at the enabling technologies level. HPC data source: [2200]. Sources for FTQC systems: PsiQuantum [2201–2204], Xanadu [2205–2207], IBM [2208], ORCA Computing [2209], Dirac [2210, 2211], Alice & Bob [2212], Quantinuum [2213], C12 [2214], and QuEra [2215]. NISQ power estimates are public vendor data besides IBM with System Two. © Olivier Ezratty, September 2026.

V/110 V wall plug with their <2 kW power drain, like a kitchen oven. Pasqal is using a similar selling point for its Fresnel quantum simulator although their future generation of quantum simulators with over 300 qubits will require a 4K cryostat to cool the ultra-vacuum pump and vacuum chamber hosting the neutral atoms with a total power footprint of about 12 to 20 kW.

Fast-forward about a decade, and we will have fault-tolerant quantum computers (FTQC) supporting thousands of logical qubits and billions of quantum gates to solve key problems like in chemical simulations. Based on the current vendors roadmaps and public data, the chart in Figure 298 shows how the related QPU would compare with the largest supercomputers deployed as of 2025. The chart contains two power thresholds: one is 40 MW, the peak power of the most power-hungry exaflop supercomputer (Aurora, at Argonne in the USA, at 38.7 MW) and two, is one 1 MW, a power that would be acceptable for two reasons. It is inline with the power of small size supercomputers and it would enable the creation of large quantum datacenters with a couple dozens of these QPUs and still be under the 40 MW threshold. This sort of setup would enable the parallelization of the multiple and identical quantum circuit shots required by the most demanding applications like QPE-based chemical simulations. If we consider that these systems will be on equal footing with their computing capabilities, this energy and cost gap will create some natural Darwinism between QPU vendors.

The notion of energy acceptability becomes of paramount importance. Quantum computers energy and power acceptability will be conditioned by the potential added value of the solutions developed thanks to them. It also depends on the electricity grid capabilities to fuel quantum computers, if their power drain is significant. As a reference, traditional nuclear reactors power capabilities range from 900 MW to 1.65 GW [2217], with upcoming SMR (Small Modular Reactor) deployments planned throughout the world with capacities ranging from 20 to 450 MW [2218]. We can also count or not with fluctuating renewable energy source capacities and their related chemical or mechanical energy storage capabilities. In the USA, the AI datacenter extravaganza is using gas turbines and diesel power engines.

The lesson from this graph is that there will be potentially a large difference in power consumption between the QPUs sold or rented. You could assert that the cheaper will win and generate a volume market. But we can't yet predict that there is a correlation between these different power estimates and the related QPU technology feasibility.

Quantum processors are noisy, with two-qubit gate error rates typically in the 0.1% to 1% range, the best trapped-ion and neutral-atom gates now sitting between 0.03% and 0.3%. This is still a prohibitive level for most algorithms and for many quantum error correction codes. In the short term, variational algorithms are developed for noisy processors in the quantum computing

paradigm called NISQ (Noisy Intermediate Scale Quantum) and using quantum error mitigation techniques. These have not yet demonstrated any clear quantum utility or advantage, at least for problems relevant to industries. In the longer term, we will rely on quantum error correction using many physical qubits assembled as logical qubits. Their number varies from a couple dozens to millions depending on the qubit technology and the target logical qubits error rates. Thanks to very low logical error rates, it will enable longer calculations with deeper algorithms.

In both cases, demonstrating a computational advantage for real quantum processors is an open field covering fundamental and applied research as well as technological developments. In some cases, quantum computers could be less energy-intensive than conventional computers to solve the same problem.

Modeling and optimizing the energetic efficiency of quantum computers must consider all classical resources used for control and error correction or mitigation. As a ratio of a performance to a resource, this energy efficiency is a hybrid quantity:

- Computational performance emerges at the fundamental quantum level, and results from the ability to control the noisy quantum processor to perform an algorithm with a certain accuracy. Understanding and optimizing these mechanisms is a matter of quantum control, quantum thermodynamics, quantum error correction, algorithms and compilers.
- Establishing satisfactory control at the quantum level requires the provision of resources at the macroscopic level, which determines the energy consumption necessary to carry out the calculation. This is the domain of enabling technologies including cryogenics, control electronics, cabling, lasers, amplifiers, detectors, classical computing resources, whose mix depends on the qubit type.

It is essential to set up a full-stack quantum computer model coupling these different levels, as well as common language and concepts.

Sizing the problem

Economics play an important role in the assessment of quantum computer energetics, with basic notions of benefit over cost ratios. Two distinct metrics are used here and should not be confused: the benefit-cost ratio, which is the economic benefit divided by the total cost of ownership and becomes attractive above 1, and the return on investment, which is the net return relative to the cost, meaning the benefit minus the cost, divided by the cost. We use the benefit-cost ratio in what follows. The energy needs of QPUs belong to their "opex" costs while their acquisition costs are their "capex". We can presume that these costs are proportional (Figure 299).

The potential economies of scale of this market will highly depend on quantum computers becoming generic

and versatile platforms, thus the need to look carefully at the various algorithm classes that could help create industry markets use cases. Also, quantum computers price and energetics will differentiate vendors in enabling or not circuit parallelization to save computing time for the most demanding applications.

Another longer-term question deserves to be asked: does the potential energetic advantage of quantum computing only depend on algorithms and applications? What will happen if quantum computing becomes widespread beyond the realms of scientific computing, and in the field of industry day-to-day operations? Are we finally going to create new sources of energy consumption that will be added to existing sources, which are already growing fast in the digital world? What will be its impact? How can it be limited? At this stage, it is too early to have a clear idea.

Answers will largely come from the emergence or not of quantum solutions for volume applications, such as autonomous vehicle routing, renewable energies routing, personalized healthcare solutions and financial services. Still, many characteristics of these applications will prevent them from being directly connected to end-user consumers. These will probably be back-office batch applications, not real time applications like social networks, ecommerce web sites, and search engines.

We can distinguish some key applications categories. Chemical and materials simulations are very promising use cases based on what we know about existing quantum algorithms but they correspond to a middle-sized market. Combinatorial optimizations to solve optimization and decision problems correspond to a much larger market potential. But there's a key algorithmic concern for multiple reasons. It seems that no quantum algorithms generating exact solutions provide an exponential speedup compared to best-in-class classical algorithms. There is also no generic proof of a practical quantum advantage with quantum heuristics and metaheuristics. Additional quantum algorithmic research is needed here so that quantum computers deliver real value in that domain [1956].

Quantum computers will be deployed and used in a different manner from typical classical cloud resources. For example, in chemical simulations, large QPUs will be required, probably with some interconnect. It seems at first glance that these interconnected systems will be 100% dedicated to a single algorithm, and for long times that can be days if not months. This will not be as flexible resources as the ones currently used in shared cloud services.

Without volume-oriented applications, quantum computers will be dedicated to niche applications equivalent to those of current supercomputers, which are mainly used in fundamental and applied research. These systems are based on calls with selection juries, not on-demand cloud access services.

On their end, volume applications will only be achievable once the quantum computing scalability will work, and millions of low-noise qubits can be operated. This scalability will probably come from fixing some of the energetic consumption issues of quantum computing. And we'll close the loop!

Then, we will have to look at the externalities of these applications and potential Jevon's effects. Namely, some new solutions will have a given quantum computing energetic cost but may help reduce the environmental footprint in other domains like in transportation. If it is well balanced, that's fine. If, on the other hand, the externalities may not be positive, like, in financial portfolio optimization tasks, you will have to think about it. It is too early to tell and is also highly dependent on the economics and affordability of FTQC quantum computers.

Another looming question is whether quantum computers will replace classical computers? If it does solve problems faster, why not? It happens to probably not be the case. First, in most situations, quantum computers will need classical computers to prepare the circuits that they run. In the case of quantum chemistry, this may become a very demanding task. In other cases, quantum processors will become accelerator of existing classical computers. The replacement thesis may work for optimization and quantum machine learning applications. For optimizations, we are still looking for quantum algorithms enabling both theoretical and practical speedups. And quantum machine learning is handicapped by classical data loading. Data loading could be efficient if it were already quantum. But most business data used in AI doesn't come from quantum sensors and is entirely classical, made of text, images and videos.

5.4.3 Quantum computing energetics

Thanks to quantum coherence, superposition and entanglement, quantum computers could showcase an exponential computing speedup compared to their classical counterparts, depending on the size and nature of the problems to be solved and on the used quantum algorithm. This computational advantage is usually predicted for ideal, error-free processors.

Energetics at the quantum level

Quantum computing energetics start at the quantum level with understanding the energy exchange mechanisms which are at play during any operations on the qubits as described in (a) from Figure 300. There are on one hand theoretically reversible operations for all unitary operations which are single and multiple qubit gates, but these operations are still driven by energy exchanges between qubit control and qubits themselves, frequently with various forms of light-matter interactions, and cavity quantum electrodynamics (cQED). Then, qubit measurements are irreversible and dissipa-




Figure 299: Measuring the impact of quantum computing can be based on assessing the ratio between its economical benefits (which are influenced by computing speed, results quality and the amount of needed data, with a balance depending on the customer requirements), and its cost structure (TCO, standing for total cost of ownership), including fixed costs (capex like for hardware acquisitions) and variable costs (opex, including energetic costs). The ratio shown here is the benefit-cost ratio, economic benefit divided by total cost of ownership, which has to exceed 1 for the investment to pay off. Return on investment is a different quantity, equal to the economic benefit minus the total cost, divided by the total cost. In some cases, the economical benefit computation is indirect when the computing solution improves non-monetary factors like customer satisfaction. It may be positively influenced by externalities like economies of scale engendered by the versatility of quantum computing platforms, which could drive prices down. On the other hand, fast obsolescence or lack of reliability could drive price up and value down. This ROI should be compared with legacy solutions and best-in-class recent classical solutions when possible. It can also help compare competing quantum computing solutions. © Olivier Ezratty, November 2025.

tive operations, and the related heat must be evacuated. Dissipation can also be engineered to mitigate the effect of noise, as is the case with cat-qubits.

Then, most qubits are operated at cryogenic temperatures. We will describe the reasons why when dealing with every qubit type in the quantum computing hardware part of this book. The main one is that these low temperatures are required to avoid the effect of thermal noise on the qubit manifold, and to preserve their coherence.

Energetics are also at play with understanding what energy must be spent to run quantum operations and what explains the gap between the energy spent at the quantum level and at the macroscopic level to generate the drive signals. Quantum noise also plays a role here. All noise sources generate losses and decoherence and their energetic underlyings must be understood. However, all the noise sources affecting qubits and qubit operations are not necessarily explained within the frameworks of quantum thermodynamics and quantum energetics. In many cases, their source are material defects, crosstalk, correlated error, control signals imperfections, photon losses in photonic components, and other related phenomena.

Another topic is the relative cost of quantum operations and their duration. There is a direct link between quantum thermodynamics and physics with the speed of the

quantum gates a quantum computer could execute, with fundamental Quantum Speed Limits (QSL) governed by energy level spacings and operating time trade-offs [475, 2219–2223] (see Figure 123). There are also trade-offs between qubit gate speed, energy consumption and some error types, particularly with superconducting qubits [2224, 2225]. Fast operations indeed have a higher energetic cost and can generate various unwanted side effects like leakage.

These various quantum energetic topics are covered by multiple disciplines, including quantum thermodynamics and quantum energetics [2226].

Quantum thermodynamics investigates energy and entropy flows at the microscopic scale, relying on concepts adapted from classical physics such as temperature, thermal baths, and heat engines that extract work from thermal gradients [2227]. This is a rich field of research [2228–2243]. It can for example help understand the role of quantum entanglement in thermodynamical systems [2244] and of quantum heat engines implemented with superconducting qubits [2245]. There are also many quantum thermodynamics and energetics concepts at play in qubits inner working and with other quantum technologies, some of them still relating to the famous Maxwell’s demon [2246, 2247]. There is even a proposal to use energy consumption as a proof of quantumness [2248].

Quantum energetics studies the flows of energy, entropy, and information between quantum systems without necessarily involving temperature or heat baths. While quantum thermodynamics implies a thermal context, quantum energetics explores engines and processes powered purely by quantum resources, such as the energetic cost of a quantum measurement and backaction [2249], the role of quantum entanglement, or of overcoming quantum noise. Quantum energetics can help determine the ultimate energy costs of qubit operations, with potential side effects on the energetic understanding of quantum error correction. Its aim is to answer questions like “*Can quantum noise be turned into an energetic resource? What is the cost of control against quantum noise? What is irreversibility in quantum processes, and how does it relate to energy waste? Can we gain new insights in complex quantum dynamics from the identification of energetic resources unlocking otherwise forbidden evolutions?*” (from Auffèves and Elouard in [2226]). Another recent example is the study of the energetics of the creation of non-Gaussian states in quantum photonics, which is a key resource to obtain a potential superpolynomial speedup with continuous variables photon quantum computing [2250].

As Harrington, Mueller and Murch wrote in a 2022 review paper “*Quantum information processing relies on precise control of non-classical states in the presence of many uncontrolled environmental degrees of freedom—requiring careful orchestration of how the relevant degrees of freedom interact with that environment. These interactions are often viewed as detrimental, as they dissipate energy and decohere quantum states. Nonetheless, when controlled, dissipation is an essential tool for manipulating quantum information: Dissipation engineering enables quantum measurement, quantum state preparation, and quantum state stabilization*” [2251, 2252]. Likewise, Matteo Carlesso and Mauro Paternostro recommend studying energetics at the fundamental quantum level to design more energy-efficient quantum devices. They promote the use of machine learning to optimize the energy exchange mechanisms in quantum computing [2253].

Reversibility. Let’s focus here on the theoretical reversibility of quantum computing operations. Classical computing is usually irreversible since it spends its time destroying information. This is a result of space-time tradeoffs made by the classical computing industry a long time ago even though Charles Bennett devised reversible classical computing techniques around 1973 [2254].

The logical reversibility of a calculation is linked to the ability to reverse it after one or more operations and recover input data from output data. This can be done at the scale of a classical logic gate or an elementary quantum gate and then up to a complete calculation. If logical reversibility is possible at the level of any gate used, then it becomes doable for a complete calculation [2255]. Today’s classical computers are logically

irreversible since they rely on two-bit logic gates that destroy information since they generate one bit with two bits, and we don’t keep the information from the two initial bits. One bit is thrown away every time. You can’t reverse a simple NAND, OR or AND logic operation. We could use reversible logic gates that do not destroy information and generate as many output bits as input bits. This would lead to a logically reversible calculation. All of this was theorized by Charles Bennett in 1973 and Tommaso Toffoli in 1980. Classical computing is a big energy spender partly because logic gates are logically irreversible.

The lower bound of energy consumption of current classical computing comes from Landauer’s famous limit of $k_B T \ln(2)$ energy dissipated per irreversible bit operation, about 3×10^{-21} J or 18 meV at room temperature, which can be the erasure of a bit or the merging of two computation paths [2256]. Even though we are far off this limit with current classical computing technologies, this lower bound could be avoided with logical reversible computing. The implementation of this logical reversibility by rewinding calculations would reduce the energetic cost of classical computing, the energy spent in the forward calculation being potentially recovered in the reverse calculation. It was not a chosen path for various reasons. First was the steadiness of Moore’s empirical law for many decades. Second, a reversible classical architecture has a significant overhead in the number of transistors used. In simple words, you need more memory.

Gate-based quantum computing is logically reversible. All gates implement a reversible unitary transformation. It can be easily visualized with single qubit gate operating in a Bloch sphere. This process is at play with the uncompute trick which uses it to reverse computation on some qubits (we’ll see that in the algorithm part of this book). This reversibility is also linked to quantum operations being linear. In the real world, qubit operations are not entirely reversible due to the effects of noise and errors.

Thermodynamic reversibility is another matter and can be obtained when the system is continuously balanced with its thermal bath. It requires handling operations in a quasi-static way, namely, slowly and with logical gates requiring a minimum energy spending. This is the field of adiabatic computing.

Gate-based quantum computing is logically reversible because it uses unitary operations which are all mathematically reversible. Qubits readout is the only logically irreversible operation when it collapses qubit states to a basis state. Measurement Based Quantum Computing, which relies mainly on measurement during the entire calculation, is irreversible by construction. This is why it is also called 1WQC for one way quantum computing.

Logical reversibility says that the map can be inverted, which works with unitary gates. It does not imply that the physical control pulse implementing that gate is

Q Counter intuitive facts

Many common wisdom about quantum computing energetic costs can be challenged. Here's a quick overview of these.

Cryogenics: all qubit types actually rely on some form of cryogenics. It is usually a variable cost vs physical qubit number. And cryogenics won't be the dominant energetic cost in most cases due to improved efficiencies. Also, an FTQC system operating at 4K may consume more energy than those running at 15 mK (more on this in 5.4.4 and 5.4.4).

Electronics and photonics: control electronics and photonics may become the dominant energetic cost in solid state platforms.

Error correction. Error syndrome detection may dominate energetic costs for some qubit modalities, particularly with silicon qubits.

Science vs engineering. There is still room for progress at the fundamental level (see 5.4.3). Enabling technologies also play a key role in enabling energy efficient FTQC. There are lots of interdependencies between stack components.

Classical computing. Its energetic costs must be considered before, during and after a quantum circuit run. Also, circuit generation and compilation may become a variable cost vs the problem data in most cases (see 5.4.3).

Economics. Energy can represent >30% of an FTQC total cost of ownership. FTQC cost could double the energetic bill of HPC centers if each center buys a FTQC system with the expected energy distribution based on vendors data (see Figure 298).

thermodynamically reversible, since real pulses dissipate in the control chain whatever the logical operations it drives. Likewise, measurement is not automatically dissipative. What carries the Landauer cost is the erasure or reset of the classical record, and of the ancilla qubits, rather than the measurement event itself. The amplification of the qubit readout signal is usually where most of the energy actually goes. A readout leaves the qubit state unchanged only when that state is already aligned with the measured basis states $|0\rangle$ and $|1\rangle$. In practice, quantum computing is not thermodynamically reversible. It would be reversible in the absence of noise and if measurements were not changing qubit's internal states. Achieving physical reversibility would also mandate that all non-quantum qubit control electronics rely on physically and thermodynamically reversible processes or at least be energy-saving operations. One way to achieve this would be to use adiabatic and reversible electronic components working from within the cryostat, but it is not possible, particularly at the DAC/ADC levels, given these analog/digital pulse signals conversions are seemingly not reversible processes. Another explored avenue is ABRC for **Asynchronous Ballistic Reversible Computing**, promoted by Michael P. Frank's team at the DoE Sandia Labs in the USA. They plan to implement it with Josephson junction circuits [2257, 2258].

So, although quantum computing is reversible from the physical and logical standpoints, in practice, it is not a path to reducing the energy consumption of quantum computing. Its potential energetic efficiency lies in its capacity to leverage interferences to efficiently find a solution to some computational problem.

Quantum Energy Teleportation (QET) is a protocol that enables the activation of local vacuum energy through the consumption of pre-existing entanglement

and classical communication. It happens only at the quantum level and can't be used to teleport macroscopic quantities of energy. QET works by having a sender perform a local measurement that injects energy into the system and then classically communicates the result to a receiver, who applies a specific operation to extract positive energy from their local region, momentarily creating a negative energy density. Proposed by Masahiro Hotta in 2009 [2259], this phenomenon respects causality and local energy conservation while enabling some quantum technologies use cases like algorithmic cooling and maintaining robust qubit entanglement in quantum networks [2260, 2261].

Quantum batteries is another burgeoning field of quantum thermodynamics and energetics that we already described in 3.10.11. Its applications to quantum computing are studied by James Quach in Australia for the improvement of the energy efficiency of superconducting qubit drive [754, 2262]. It has already been experimented at small scale. Scaling faces some challenges like frequency crowding, battery decoherence, the recharge problem, control precision, and crosstalk but this can provide interesting ideas to researchers in the field of superconducting qubits.

Energetics at the macroscopic level

Operations on qubits happen at the quantum level but they always start at the macroscopic level. Macroscopic electronics or photonic signals are generated classically and then conveyed to the qubits. Estimating energetic resources for quantum computing requires a full-stack approach encompassing all classical electronics, photonics and cryogenic costs, plus all classical computing costs (preprocessing, processing, error correction) and the number of circuit shots. All energetic costs

are of classical nature. They depend on the algorithm structure, size and goals like in accuracy for a chemical simulation, and, on the noise affecting qubits at the quantum level. Then, comparisons can be made between different quantum computers and between these, and entirely classical solutions using best-in-class classical algorithms. Figure 300 inventories all the classical computing, classical electronics and photonics, and other enabling technologies to account for.

Let's break down the power consumption of a typical quantum computer:

Control electronics power consumption varies greatly from one technology to another and depends on the number of physical qubits managed, which will be counted in millions with large scale quantum computers (FTQC).

Superconducting qubits are driven with microwaves produced outside the cryostat with electronics coming from Zurich Instruments, Qblox, Quantum Machines, Keysight and the like. Microwave readout is costly in bandwidth, requiring Gbits/s of data streams per qubit.

Pulsed microwaves can be generated by cryo-CMOS chips sitting in the cryostat at the 4K stage. It can help reduce the wiring clutter entering the cryostat but is not necessarily reducing the energetic footprint since it increases the cooling budget requirements. It seems to fit the needs of intermediate systems with only a couple hundred or thousand superconducting or silicon spin qubits. Another option is to use superconducting electronics (SFQ) like the ones developed by SEEQC but it still requires some work to avoid side effects like back-action on the qubits.

Classical control electronics power can gain at least several orders of magnitude. IBM plans to go from 100W to less than 100 mW per qubit [2263].

Trapped ions qubits control is performed with lasers and/or conventionally generated electromagnetic pulses. Its cost scales with the number of qubits.

Cold atoms qubits control exploits a couple of lasers and a mix of SLM and AOD systems that potentially support a thousand qubits with modest power consumption. But laser power can grow significantly as the qubit number exceeds a couple hundred.

Photon qubits power drain seems more important for photon detection (about 7.5W per qubit) than for photon generation (about 1mW per qubit, source: Qunela). Power consumption also comes from the control electronics driving the nanophotonic circuits (phase controls, ...). Superconducting based photon detectors are more demanding with cooling.

Vacuum is required for superconducting and silicon spin qubits, while cold atoms and trapped ions qubits need ultra-high vacuum. Photons do not need it. With superconducting and silicon qubits, vacuum comes from pumps and dilution refrigeration cooling. Cold atoms require only 100W for the ultra-vacuum pump plus about 300W for its cooling at 4K. Trapped ions systems use

heating strips covering the vacuum chamber with a process that can take weeks. As with cold atoms, ion traps chambers may be cooled at 4K to avoid thermal photons entering the system and to improve UHV quality.

Cryogenics consumes up to 16 kW (for about 50 qubits) and, for an IBM Quantum System Two with a Bluefors KIDE and 9 pulse tubes, an estimated 105 to 110 kW. That second figure is an engineering estimate built from the bill of materials of the cryogenic platform and not a an official measurement. You must add here the energetic cost of cooling the water circulating in the cryostat compressor. Cryogenics will be required for cold atoms. Atoms are cooled with laser beams and SLM/AOD driven tweezers and under ultra-high vacuum but their vacuum chamber is now cooled at about 4K to maintain a good quality vacuum which is generated by room temperature ionic-UHV pumps. In most systems, the cryogenics consumption is usually continuous, without much variations between thermalization and operations. Thermalization lasts about 24 hours for dilution refrigerator systems used with superconducting and electron spins qubits. The cryogenic cost has different scale properties depending on the qubit types.

With superconducting qubits, energy costs scale with the number of physical qubits but also depends on the way they are controlled, the nature of control signals (microwave pulses, DC pulses) and how control signals are multiplexed and demultiplexed. One must not forget to account for the cost of water cooling the cryostat compressors. More powerful cryostats can be created with more pulse tubes and dilutions with a gain of an order of magnitude for the available cooling power. Other optimizations can be implemented to increase the available cooling power at very low temperatures by getting closer to the theoretical Carnot efficiency, or with balancing the thermal load across the various dilution cold plates [2264]. It seems possible with large cryostats to use more cooling stages. There are still significant scale constraints for FTQC QPUs which will require millions of physical qubits.

There are some known engineering paths to increase both the cooling power and cooling yield of 15 mK and 1K cryostats. IBM's Starling (200 logical qubits and 100 million gates by 2029) and Blue Jay (2,000 logical qubits and 1 billion gates by 2033) are planned to spend about 20% of their energy cost on cooling and 80% on control electronics due to better energy efficiency of centralized liquid helium production.

Computer control is used with all types of qubits. They all require from one to three control servers that drive the qubit gates and readout devices by exploiting compiled quantum software, that transforms qubit gates into low-level instructions for qubits initialization, control and readout. We usually have one server for compiling the quantum code, another one for running the compiled code and handle pre- and post-processing, and a third one for communication with the outside world, and potentially, quantum circuits batch manage-




Figure 300: Full-stack functional decomposition of the energetic cost of a quantum computing solution, all of these having a direct or indirect energetic cost and are supported by some hardware component. The arrows indicate how functions and function groups influence other functions. (a) starts by generating a quantum circuit, compiling and optimizing it so that it can be executed in the quantum processor. Other software functions play a role during the execution of the circuit like quantum error syndrome detection, or a classical optimizers with variational NISQ circuits. Post-processing uses the output of multiple circuit runs to derive a value of interest. All these software functions are supported by existing classical computing hardware with known energetic specifications, from simple laptops to GPU racks depending on the workload. (b) corresponds to all active and passive electronics and photonic components driving the physical qubit gates and readout operations specified in the compiled circuit. Light orange flags active components which have their own power consumption. Their thermal dissipation usually requires traditional air-air or water-air cooling. Dark orange boxes flags thermal dissipation from active or passive control and readout devices operating at cryogenic temperatures, which influences the cost of cryogenics. Single photon sources and detectors operate at the quantum level, thus their dark grey background. (c) corresponds roughly to what happens at the quantum level. Depending on the compiled circuit size, classical control drive the coherent gate operations which are reversible on paper and other operations like readout. These operations, including the identity operations (idle state) are affected by multiple sources of thermal and non-thermal noise which conditions the overhead of quantum error correction and therefore, the cost of all classical control costs, thus the arrow towards the right. Some quantum level operations like qubit readout are dissipative but, being at the quantum level, have a very small impact on the cryogenic costs in comparison with control cryogenic costs. (d) cooling at room and cryogenic temperatures are more or less proportional to the number of physical qubits in most qubit modalities. All cryogenics require some form of water-air room temperature cooling. (e) vacuum and ultra vacuum pumps are to be accounted for although their energetic bill is relatively small. © Olivier Ezratty, 2023-September 2026.

ment. These servers are networked, either on premises or in the cloud and via conventional network switches. They represent a limited fixed cost with an estimated consumption of between 300W and 1 kW. Part of the control computing could be moved into the cryostat for superconducting and electron spin qubits, to implement autonomous error correction codes. The control computer would then only drive logical qubits and not the physical qubits of the configuration.

Error correction conditions the power consumption of a fault-tolerant quantum computer. It is related to the classical electronics and computing resources used for qubit control (AWGs, mixers, DACs), qubit readout (mixers, ADCs, FPGAs), and error syndrome decoding (optimal maximum-likelihood decoding of a general

stabilizer code is NP-hard, even #P-hard for degenerate codes, although polynomial-time heuristics such as MWPM, union-find or belief propagation are used in practice for surface and qLDPC codes, with a current poor understanding of their classical computing cost and energetic footprint) and correction. One key parameter is the ratio between the number of physical qubits and logical qubits, which depends on physical qubit fidelities. The higher this one is, the lower the ratio of physical/logical qubits. It also depends on the algorithm size and its target performance. In the FTQC regime, the number of qubits to control will be multiplied and generate high energy consumption. However, error correction codes may run up against another wall: the scale dependence of qubit noise with some qubit modalities.

Namely, qubit gates and readout fidelities may decrease with their number. This may have the consequence of reversing the effect of increasing the number of qubits used in error correction. The error rate of logical qubits gates then increases, instead of decreasing [2265]. The same problem arises with surface codes although their overhead seems lower than with concatenated codes.

Required physical error rates are code- and noise-model-dependent and there is no universal fidelity threshold, the target depending on the code, the operation mix, leakage, erasure bias, correlations, the decoder and the target logical error rate. However, the 99.9% figure is routinely used in most vendor resource estimates, meaning an error rate below 10^{-3} for all operations (gates, readout, idle). It is not observed at scale yet, the exception being trapped ions on small registers below 100 physical qubits. In their work modeling a full-stack energetic cost of a superconducting qubits computer, Fellous-Asiani et al. found out that, from the energetic footprint standpoint, it is way more efficient to use room temperature electronics than cryo-CMOS due to the overhead cost of their cooling. This result moves the scalability burden cost on the wiring and its multiplexing. On top of that, control electronics have an energetic bill that is much bigger than the cryogeny used for the electronic components sitting in the cryostat (cables, filters, attenuators, qubit chips, circulators, amplifiers).

On syndrome decoding, companies like Riverlane plan to implement the classically costly QEC syndrome decoding with ASIC circuits instead of FPGA. This will considerably reduce the cost of the classical part of QEC. Their Collision Clustering decoder gives the first published order of magnitude: an ASIC implementation keeping up with MHz-rate syndrome generation for a 1,057-qubit surface code occupies 0.06 mm² and draws only 8 mW, against a few watts for the equivalent FPGA [2266]. Riverlane also announced in 2024 its endorsement of the Quantum Energy Initiative.

Recent development of non-local error correction codes (qLDPC, bivariate bicycles) that have a lower qubit count look promising. They rely on better qubit connectivity, which improved in the case of cold atoms with the QuEra/Harvard/MIT December 2023 experiment with 48 logical qubits and movable atoms [1856]. IBM's roadmap also now includes a mid-range in-circuit coupling enabling qLDPC and a lower estimate of the number of physical per logical qubit.

Many of these quantum computer components have a variable energy cost depending on the number of qubits, including the cryogenic side. Indeed, the electronics embedded in cryostats release heat in approximate proportion to the number of physical qubits used. This heat must be evacuated within the cryostat. The consumption of control electronics also generally depends on the number of qubits. It seems that, up to a thousand qubits, control electronics is a fixed cost for cold atoms. Only vacuum creation and the control computer seem to be fixed costs.

Engineering challenges

Some significant engineering is required to optimize a multi-parameter system:

- **Classical costs.** All algorithms resource estimations should be studied in a full-stack manner, including all classical computing preparation costs and compilation costs, including all quantum circuit generation costs, particularly for business operations applications that need several runs per day. Classical energetic costs must be embedded in any quantum computing energetic costs estimates and benchmarking recommendations. It can lead to various trade-offs, on the way to allocate the workload across the classical and quantum part of the solution.
- **Compilation costs** are a new area of investigation for large-scale FTQC applications. They appear to be a variable cost relative to user data, contrarily to the classical von Neumann model. It may handicap so-called “real-time” or near real-time use cases, particularly related to solving combinatorial optimization problems. As a result, analyzing compilation costs and its energetic footprint appears to be a new area of research. Also, early work is done on optimizing the energy efficiency of compiled circuits [2267].
- **Error detection.** Likewise, the cost of the classical part of quantum error correction, including syndrome error detection for large logical qubits, is to be studied from a theoretical and energetic standpoint. Early hardware work is starting to report area and power budgets alongside decoding latency, which is the right framing [1820, 2266].
- **Computing time.** Another critical element is the speed at which quantum circuits are executed, considering useful algorithms like quantum chemistry may have very long computing times for large molecules with over a hundred orbitals, in the days-months and even years and decades range. In a clear quantum advantage regime, computing times may become part of a solution's cost and liability as well as benefits. As a result, creating faster qubit gates, readouts and error correction operations will be critical to mitigate that cost. Three complementary ways can speed up computing operations: reducing the overall software cost (corresponding to the total energy cost), increasing qubit gates, readout and error correction speed, and parallelizing circuit shots over multiple QPUs (provided that these QPUs or cores should have moderate power drain and economic cost, circling back on the economics of quantum computers [2268]).
- **Models.** Optimizing the energetics of quantum systems is a complicated problem based on multiple and sometimes contradictory constraints and figures of merit. The first proposed model was MNR (“Metric Noise Resource”), which was created in 2021, which

helps assess the energetic performance of quantum technologies. It deals with finding the right trade-offs between noise and resources to fight it, to reach some algorithmic metric [2269]. It could be expanded with ad hoc models and methodologies encompassing use cases, various economic concerns and all computing stacks. Mathematical models for broader optimizations could be created. Also, the role of algorithms and software engineering innovations seems critical to optimize quantum computing energetics beyond the MNR model which focuses on the physics parts of computing.

- **Scientific unknowns.** Outside the energy domains sit many scientific unknowns and unknown unknowns, particularly experimental ones. What is the shape and form of quantum noise at large scale? How can large-scale entanglement operate with millions of physical qubits? Will we discover new forms of correlated errors in large quantum systems? Most of these unknowns may have an impact on the overhead of quantum error correction and therefore on the energetics of FTQC QPUs.
- **Classical computing progress.** Quantum computing is developing in a field where classical computing is also making much progress, in terms of hardware, algorithms and software, and even power consumption. This must be accounted for when comparing the energetic footprints of classical and quantum solutions.
- **Distributed quantum computing.** Scaling quantum computers to support several thousand logical qubits with the best error correction codes will need quantum interconnect solutions, to efficiently and quantumly connect many quantum processors or cores. Most qubit modalities seem to be unable to escape that route. As a result, estimating and optimizing the power cost of these interconnect platforms will be critical, including their accrued cost for code compilation and circuit execution. Vendors who can avoid this route while supporting thousands of logical qubits will have a significant economic and energetic lead. The temptation is great to create ever larger quantum computers, with giant cryostats in the case of superconducting qubits, like we see with IBM and Google’s roadmaps. Another approach would be to create distributed architectures of quantum computers linked together by quantum connection based on entangled photons, a choice made by IonQ with their trapped ions qubits, noticeably because it is difficult to scale these qubits beyond a couple dozens. In theory, this would make it possible to create computing clusters that, seen from the outside, would create a single computer, a bit like a large classical server cluster. This will be conditioned by the capability of converting qubits states to photons qubits states and by the resulting qubit connectivity between the various quantum processors units of this quantum cluster.

But this is not just about “connecting” qubits. Interconnect architectures are about creating fragile entangled states between qubits using the intermediary of photons. These may create some statistical overhead on their own, which must be boiled in for both assessing the real obtained quantum computing scalability and the related energetic footprint.

Pessimistic assumptions

Several preprints and peer-reviewed papers have been published since 2022 with existing or future quantum computer energetic consumption estimations. They started to shed a light on the topic. But for multiple reasons, one of which being the newness of the approach, their underlying methodology often missed various sources of power drain, and they were drawing optimistic conclusions too rapidly. Others were too pessimistic because they didn’t account for the foreseeable improvements in many quantum computer hardware and software components.

Outdated assumptions. The Rand Corporation produced an estimation of the energetic cost of a QPU able to break RSA-2048 codes in 2023 by consolidating various resource estimations. It estimated a rough, power cost of 6.25 W per physical qubit (with some mistakes) and simplistically multiplied by a number of physical qubits (about 20 million) but reminded readers of the many technology uncertainties about scaling [2270]. The paper did not really account for some potential future progress in electronics design, cryo-electronics, error correction code improvements and the like. In a word, it shows that such analysis can be quickly outdated by scientific and technology progress.

Narrow scope. A study published in Nature Reviews Physics in 2025 [2271, 2272] found that quantum computing was 5 orders of magnitude more expensive in energy than classical computing. The example was based on a single 5-qubit circuit without any scaling law. It is a known fact that under 20 to 30 qubits, classical emulation is usually cheaper than a quantum computer, let alone a best-in-class classical algorithm. Another paper aimed to optimize the trade-offs for running a VQE between the output accuracy and the number of circuit shots, but in a useless regime using only 5 qubits where classical always beats quantum computing, and with no indication of scaling laws [2273].

Quantum optimization. An early 2024 work on quantum algorithms energy optimization focused on comparing NISQ quantum optimization algorithms between gate-based QAOA and a D-Wave annealer used natively. The conclusion was that QAOA is worse than D-Wave native annealing which is worse than a regular CPU in computing time and energy spent. And they even didn’t account for the classical optimization for the variational QAOA circuit [2274]. This exemplifies the overarching challenge for the quantum computing science and industry ecosystems: find quantum algorithms which bring a quantum advantage for solving optimization problems.

Worrying numbers. A 2026 review coauthored by DoE ORNL scientists including David McCollum estimated that quantum computers could someday consume up to 80 GW in a stretch high-penetration scenario, and just in the USA [2275].




Figure 301: Under a deliberately high-penetration US scenario, FTQC systems would need 80 GW of power, and converting that US scenario into a worldwide helium 4 comparison adds further assumptions of its own, starting with the geographic scope. The assumptions are rules of thumb based on 3,000 FTQC systems being used, each consuming an average 27 MW, based on some extrapolation related to superconducting qubits and also, broad scale usage of these systems across all the promised domains: scientific computing, optimization problems and quantum machine learning. Consuming such energy would require quantum computers to bring a lot of value to the industry, way beyond what is expected today. Source: [2275].

It also looked at water, helium 3 and 4, and nitrogen consumption. That high penetration scenario would require 120 million liters of liquid helium 4 and 26 million liters of gaseous helium 3. It turns into about 90 million cubic meters of helium 4 gas considering a compression ratio of 1 to 757. As of 2026, worldwide helium consumption was at around 160 million cubic meters [2276, 2277] (Figure 301). At the worldwide level, a stretch

scenario would have quantum computers consume all helium 4, at the expense of MRI and other industry usages.

These 80 GW, which cover a US scenario and not a global one, correspond to more than all AI GW consumption planned worldwide for 2030-2033. So, where's the catch? It shows up at multiple levels in such estimations. One is the number of FTQC units, at 3,000, which depends a lot on these systems ability to solve industry-relevant problems.

If FTQC is mainly useful only for scientific computing, the installed base will be much smaller. Second, the assumption is that the average power consumption would be 27 MW, which is the power class of the handful of exaflops classical HPCs available today. Who could afford deploying 3,000 systems of that size and cost given the existing distribution of HPCs? Third, 27 MW sits in the high-range of expected FTQC systems power consumption. We can expect that some qubit modalities and technology developments will enable the creation of large-scale FTQC systems with less than a MW of power consumption. The low penetration scenario in the paper is more reasonable with 300 FTQC systems being deployed. With large scale optimized FTQC systems consuming each 4 MW, it gives a total power of 1.2 GW, only for the USA. You could triple it for a worldwide estimate, landing at 3.6 GW.

Optimistic assumptions

This corresponds to papers, mostly from the academic world, which describe some form of quantum energetic advantage. A majority of these papers share various shortcomings with overly optimistic assumptions (like with physical qubit error rates, gate speeds, qubit control power requirements, etc.), missing many costs (control electronics, classical control, and classical computing costs), being based on other flawed papers, being too theoretical (like using some Landauer limit for some energetic cost estimate, lack of data and algorithmic resource estimates) and wrongly estimating equivalent classical energetic costs. A last one, is a disconnect from industry vendors technology plans, particularly related to enabling technologies and QEC/FTQC, while these plans are also generally too optimistic on qubit operations fidelities. This illustrates the need for assembling research teams with a broad set of skills.

Full-stack pioneering. This approach is mandated for seriously estimating the energy consumption of a whole system and also, to find ways to optimize it. In 2022, Marco Fellous-Asiani et al. built a full-stack model of a superconducting qubit quantum computer with the aim to find parameter regimes which minimized energy consumption for a given computing performance [2269, 2278]. It considered typical algorithms used for optimization, physical simulations, quantum machine learning, and cryptanalysis for integer factorization. The full-stack modeling integrated the sources of quantum noise affecting qubits, the conventional qubit control

resources, then the sources of heat dissipation involved in the whole material chain and in particular in the cryostat. It also accounted for the QEC code size, initially a (heavy) concatenated Steane code, and then a surface code. The model established a relationship between microscopic processor parameters such as qubit fidelity, with macroscopic qubit control parameters. It was an interesting basis to optimize the energy consumption of the whole system, under the constraint of reaching a targeted computational performance.

The paper found that an order of magnitude gain in qubit fidelities (which is quite hard to obtain) could lead to an energy consumption saving of 100. For control electronics with a highly optimistic assumed power consumption of 2 mW per qubit [2279], it seems preferable to run it at room temperature than at cryogenic temperature, providing cabling multiplexing solutions can be implemented. The technological constraint then lies in the wiring, which must be simplified, essentially by using advanced (and future) multiplexing techniques. Another option would be to use superconducting electronics running at the processor level or at the 4K stage, like the ones developed by SEEQC that we describe later in the enabling technologies part of this book.

With the proposed model, the possibility of an energy-based quantum advantage was investigated. It computed the minimum energy consumption of a fault-tolerant quantum computer to factor an N -bit integer and compared it to classical solutions [2280]. It compared NISQ systems and their classical emulation equivalent, but not best-in-class classical algorithms equivalents. The last known classical record was obtained in 2021 by an Inria team, on a Germany-based supercomputer, for factoring an 829-bit key [2281] with a consumption of 965 GJ, corresponding to 1.3 MW of power spent over 8.6 days. The model showed that a quantum computer operating with qubits 2,000 times more faithful than Google Sycamore (quite a feat), combined with Steane's code would require 2.7 GJ, corresponding to 2.9 MW for 16 minutes, to factor such a key. That would be 350 times less energy. Breaking a 2048-bit RSA key on a QPU of the same type as above, the energy consumption would be 38 GJ, corresponding to 7 MW for 1.5 hours. Estimations were made for different key sizes in the classical and quantum cases, giving access to some energy efficiency in each case. An energetic quantum advantage would start to show up with $N=848$ for the key size. It would be different in nature from the computational advantage, which considers only computation times [2282, 2283]. Both advantages would thus be achieved for different key sizes. Still, the proposed error correction code was resource-intensive, and its cost would be much lower with, for example, a surface code, or potentially, a qLDPC code with better qubit connectivity.

So, are we sure to get this sort of energetic advantage before any computational advantage [2284] ? Not yet. The model was highly theoretical and with many highly optimistic assumptions on qubit operations fidelities, signals

multiplexing, control electronics power, and cryogenic yield. It was based on using a monolithic superconducting chip while these may not scale beyond a couple hundred qubits and require some microwave or infrared entangled photon interconnectivity. This technology will add its own burden, lowering two-qubit gate fidelities and increasing energy consumption. The paper also could not account for the classical computing cost of error correction. It may be very expensive and even become the main energy consumption source in some systems.

The interest of the proposed model was to highlight technology interdependencies, including qubit operating temperature, which can help make sound choices in quantum computer design. Its use case based on Shor integer factoring is not far from many chemical simulations resources estimates that sit in the thousands of logical qubits. It deserves to be applied to various quantum computing architectures, qubit types and paradigms. The model was later applied to a quantum kernel algorithm, showing that there was no looming quantum computing advantage as well as any quantum energetic one [2285].

Missing energetic costs. In one example related to QPU energetic cost estimates for Bitcoin mining, the energetic consumption was based on Landauer's (tiny) lower bound applied to the cost per qubit gate, and is not considering any qubit control, cryogenics and classical error correction energetic costs [2286, 2287]. The advertised result touting some quantum computing energetic advantage for Bitcoin mining is then misleading. Another 2024 paper did showcase some energetic advantage for Bitcoin mining, that was focused on SHA-256 hashing while it is known that proof-of-work in Bitcoin mining requires a Grover search algorithm that doesn't scale well [2288]. In the longer term, the startup BTQ (Taiwan, Canada) plans to use some photonic quantum computing to implement proof-of-work and save energy for bitcoin mining and Blockchain processing [2289]. This seems less far-fetched although still futuristic.

Missing oracle. In another example, Meier and Yamasaki demonstrate that “*quantum computation achieves an exponential energy-consumption advantage over classical computation for Simon's problem*”. The estimate is based on a single (Simon) oracle based algorithm without specifying how the oracle is implemented nor its cost, which may have a significant computing cost. It doesn't consider the cost of error correction or QPU interconnect. It also determines lower and upper energy consumption bounds formula. It shows an asymptotic exponential gain in energy consumption but does not account for the classical cost of quantum error correction [2290]. The work has since been published in PRX Energy.

Confusing power and energy. A 2023 preprint from Junyu Liu et al. created some economic laws showing the energy economic advantage of quantum computers but seems off the mark when estimating the energetic

costs of these systems [2291]. For example, the authors wrote that “*it has been reported that Google’s quantum devices require approximately 15kW for the complete experiment. Since this energy demand is primarily tied to cooling, it does not significantly fluctuate with the number of qubits*” which contains three factual errors. First, 15 kW is a power metric (in W), not an energy consumption (in Wh). Energy consumption depends on the total classical and quantum computing time. Second, Sycamore’s experiment had a power drain of 26 kW as mentioned page 51 in the Google paper supplemental materials [2292]. Third, energetic costs scale with the number of physical qubits! Then, they make some assumptions on the energetic cost of gate-based Rydberg atoms quantum computing using kJ per two-qubit gates. But they didn’t consider the number of shots, error correction costs or atoms preparation costs. Why bother? Well, this paper was mentioned in an April 2026 McKinsey report touting some quantum energetic advantage [2293] (slide 95).

Confusing NISQ and FTQC regimes. A May 2026 multi-team survey led by Miquel Carrasco-Codina investigated the energy consumption and potential energetic advantage of NISQ systems [2294]. It had interesting insights like considering that there was not much difference between fixed and variable costs related to idle versus active times in control electronics. It led them to compute energetic costs based on circuit execution duration, and they addressed computing time alongside energy efficiency, which is the right framing. However, when looking at energetic costs, they considered circuit depth ranges well beyond the accessible NISQ regime (their Figure 5, where roughly 98% of the x-axis corresponds to circuit depths not reachable with NISQ devices). The paper also needs to refine its terminology around algorithms, circuits, shots and computations and highlight that, for example with a VQE, the number of circuit shots is determined by the algorithm and its required accuracy, scaling as $1/\epsilon^2$, ϵ being the expected chemical accuracy, and on the number of Pauli strings, ansatz parameters, and the efficiency of the classical optimizer.

Single or multiple shots accounting? A 2024 paper from Yasser Omar et al. did show that a quantum energetic advantage could be obtained with trapped-ions QPUs starting with 43 qubits [2295], using the case of a QFT compared to a classical DFT. It was based on a single probabilistic circuit execution of the QFT when it requires multiple shots to generate a full frequency-domain decomposition like a classical DFT which yields a set of complex numbers with this decomposition. I reached out to the authors and they acknowledged the problem. They also did extrapolate this result to the implementation of a Shor integer factoring case when its most efficient version is now based on a low-cost semiclassical QFT, the burden of Shor’s factoring algorithm being in its exponentiation (or period finding) part, which is a more complicated circuit than a QFT. Finally, computing times were not estimated in the pa-

per with large QFT in the quantum energetic advantage regime above 43 qubits. The work could be complemented by resource estimates related to quantum error correction which may be required when reaching that scale or at least above 100 workable qubits.

Abusively theoretical. This is the case with this 2025 paper touting some “*Energetic advantages for quantum agents in online execution of complex strategies*” [2296]. It happens to be entirely theoretical, based on computing some lower bound, with no hardware mentioned for either a classical or quantum implementation. Mentioning some “*energetic advantage*” without any engineering is stretch. By the way, this is using only 4 qubits for which no quantum advantage is possible, whatever the qubit type. A similar entirely theoretical approach was adopted in a paper showing some quantum advantage with a quantum agent network.

Russian-doll errors. Another 2025 paper “*Towards identifying possible fault-tolerant advantage of quantum linear system algorithms in terms of space, time and energy*” [2297] illustrates several methodological gaps. It searches the conditions to obtain some energetic advantage with running an HHL algorithm. The assumptions are: qubit error rates of 10^{-5} , which is at least two orders of magnitude below what is measured today, excluding memory and data transfer costs (which for HHL is known to be costly), based on the belief that for superconducting devices, “*the primary energy costs comes from cryogenics*”, whereas full-stack models attribute the dominant share to control electronics, and a comparison made with a laptop, which is not the right classical baseline: an Apple M-series laptop delivers a few GFLOPS/W in FP64 at the whole-system level while the best HPC sit around 60 GFLOPS/W, and in FP16 an HPC node or an Nvidia DGX is about an order of magnitude more efficient. The paper makes references to other papers which are themselves muddying the water. They use reference [2270] which mentions that IBM Quantum System Two “*houses 4,158 physical qubits and operates at a power of 26kW, resulting in a power efficiency of 6.25 W per qubit*” (this QPU could host this number of qubits, but its cryogenics alone probably costs about 110kW and you need to add control electronics which cost between 10 W and 100 W per qubit, even though in the future it could decrease below 1 W). They also say that they would need some QRAM circuits for data loading but that they are classically simulable. At last, this energetic advantage would come with $O(10^5)$ physical qubits (>100K), which don’t fit in a single fridge or chip.

Missing data. This is the case with the paper *Accelerating the drive towards energy-efficient generative AI with quantum computing algorithms* [2298]. It doesn’t provide any sizing of the classical AI consumption and the related use cases (LLMs? other applications?). It then makes an inventory of various other quantum machine learning algorithms to process some parts of generative AI processes. However, none of the referenced

papers contain resource estimates corresponding to the size of existing LLMs or other AI applications. One referenced paper is about optimizing an existing neural network with a D-Wave annealer. It uses an MNIST (handwriting recognition test with 28×28 pixel images) benchmark and it happens a digital (classical) annealer is more efficient. But a MNIST neural network requires about 140K parameters while an LLM starts with several tens of billions of parameters. The paper also asserts that: “...leading quantum computing architectures also generates significant power demands, which, depending on the exact technology, may have an order of magnitude of tens of kW. Comparing that with the MW power that classical supercomputers often require suggests that quantum computers can be competitive in terms of energy consumption-related quantum advantages”. This is like comparing apples and oranges. You cannot compare an existing NISQ QPU with a supercomputer for tasks that will probably require an FTQC QPU with a much larger energetic footprint, which may sit in the MW zone for superconducting qubits per IBM’s plans!

NISQ assessments. Some optimistic assumptions are made when looking at NISQ QPU power consumption. Currently, most NISQ solutions can be emulated classically and more energetically efficiently (VQE, QAOA, QML). NISQ solutions which may bring some energetic advantages are in the analog quantum simulation realm, with Pasqal and QuEra QPUs, and potentially, with the quantum utility regime promoted by IBM with >100 physical qubits of good quality and by using efficient quantum error mitigation techniques.

Applications scope. Most of the above mentioned papers usually deal with a specific quantum algorithm and sometimes, the classical equivalents used for the comparisons are not the best-in-class or equivalent. The algorithms’ choice can generate biased conclusions in such comparisons. A quantum computing energetic advantage should be as generic as possible, like for a computational advantage.

Benchmarking quantum computing is still a debated topic between the ones who prefer low level benchmarking tools (qubit fidelities, quantum volumes, CLOPS, layer gate errors, etc.) and higher level ones based on typical case studies.

The best generic computational and energetic advantage should preferably be assessed by solving combinatorial problems like a MaxCut, chemical simulations with a Quantum Phase Estimation, given its cost lies in the exponentiation and not with its inverse QFT part, solving linear equations with HHL, and potentially, some generic machine learning use cases (which are harder to pick).

Other preprints are looking at some detailed aspects like an energy efficient way to implement qLDPC codes using analog classical control circuits [2299], with studying the lower bounds for the classical control of the qubits and

the trade-offs between its energetics and speed [2300] and, on how to optimize cryogenics [2301].

Lastly, another paper explores the interesting field of quantum code classical emulation, showing that emulating 43 qubits for 1.64 hours costs about 570 kg of CO₂ [2302]. The first full carbon footprint comparisons between supercomputers and quantum computers have been done, but with no real use case benchmark [2303]. It provided an interesting comparison between the largest US supercomputers and a 100-logical qubit QPU based on Nord Quantique’s GKP qubits, which are still to be demonstrated at scale, and with no real benchmarked problem other than mentioning Google’s supremacy experiments which solves no problem given they are using random circuits and no input data.

5.4.4 Qubit type differentiation

Each qubit technology comes with its own assets and challenges regarding their energy consumption.

Superconducting qubits

This is the qubit type with the most abundant literature on energy efficiency and optimization.

Spontaneous microwave emission by the qubits is a dissipative process engendering errors and decoherence. Energy exchanges happen between the qubit and its controlling microwave during a qubit gate operation [2304]. Even qubit dephasing drives energy dissipation [2305, 2306]. Error mitigation can make use (among various other techniques) of dissipation engineering with energy baths whether it is handled with bosonic qubits like cat-qubits or with programmed error correction. Engineering dissipation can also help efficiently prepare (entangled) Bell states with two qubits. Measurement energetics is also studied with the Zeno effect with measurement backactions, how to optimize measurement operations with various types of microwave light (single photon, coherent light, thermal light) [2307–2309], the connection between quantum measurement and error correction [2310] and ways to purify the state of a single qubit with a quantum thermodynamic method [2311]. There are also specific cooling mechanisms for superconducting qubits and even some connections between qubit thermodynamics and the way to optimize computing at the compiler level. Finally, in the internal debates between the types of superconducting qubits, let’s note that fluxonium qubits may have different energy consumption since being driven by lower frequency microwaves.

Cat-qubits from Alice & Bob did show the capability to correct flip-errors at the cost of a moderate growth of phase error, which can be corrected with 1D qLDPC codes, at least until reaching a logical qubits error rate of 10^{-8} . Alice & Bob plans to implement a 100-logical qubit QPU with a single chip containing 2,000 physical qubits with a power cost of 160 kW.

Q Quantum computing energetic estimations guidance

This part lists several shortcomings from the few academic papers who try to estimate the energetic cost of quantum computers or describe some energetic advantage. Here is a resulting summary providing some guidance on how to produce better quality energy consumption estimates.

- Describe your **assumptions**, their potential limitations and how they could influence the quantitative outcome of the research work.
- Explain the difference between **theoretical and practical** energy consumption estimates. For example, include realistic costs for qubit control and not just stick to some Landauer limit.
- Understand well how **real-life quantum algorithms** work. Account for the total number of gates and circuit shots, particularly with NISQ variational algorithms.
- Include **all computing costs**, including data loading quantum circuits, all classical costs to prepare quantum circuits and post-process their outcome, costs of error correction or mitigation, and for large circuits, the compilation and opti-

mization costs. Sometimes, the data is not available, but still mention that these cost need to be accounted for.

- Account for the overhead costs of **quantum error correction** to run a circuit having a size larger than what can be supported on a NISQ quantum computer, beyond 5K to 15K gates and idle operations.
- Be careful when **using other references** which are imprecise on the existing costs of NISQ QPU. Don't extrapolate it too fast to larger size systems.
- You can use **assumptions on future capabilities** of qubits (particularly fidelities) and enabling technologies (cryogenics efficiency, control electronics costs) but make sure they are reasonable and also well timed. Don't extrapolate on the existing costs of control since various technologies are also bound to reduce their costs like moving from FPGAs to ASICs for some functions.
- Consider the need for and overhead of **QPU interconnect technologies** that will be needed for some qubit modalities when exceeding a certain number of physical qubits.

Similar interesting results came from Nord Quantique with their first GKP superconducting qubits with unbiased noise. These are good news, if and when that works. At scale, these systems could potentially support 100 logical qubits with a really moderate power drain.

D-Wave's quantum annealer is also using superconducting qubits. The company often puts forward the quantum energetic advantage of its systems in its communication, even with AI use cases, which are unrelated to LLMs [2312]. It is clouded by shady comparisons using their hybrid solver which is mainly a classical computing tool, and comparing quantum heuristics and classical exact solvers. One remaining interesting case is the Ising model physics simulations implemented on D-Wave systems, but they are also contested by quantum inspired solutions. Some academic work is published on the energetics of these systems [2313, 2314]. D-Wave is also promoting the use of its annealers to improve the energy efficiency of crypto tokens proof of work [2315].

Regarding cryogenics, superconducting qubits are the ones requiring the lowest cryogenic temperature, between 10 and 20 mK. Its cost is larger than those requiring 2K to 4K cryogenics like photonic qubits. But the balance may be favorable to superconducting qubits, which do not suffer from the inefficiencies of some photon sources.

Silicon spin

Their energetics are also studied. Their operating parameters and controls are a bit different from superconducting qubits, with a richer mix of microwave pulses and direct current and operations at a potentially higher temperature in the 1K range. Some quantum energetic advantage can even be found at the scale of a *classical* full-adder implemented with three silicon quantum dots, with a gain of three orders of magnitude [2316] and with entanglement generation between electron spin and photons in devices mixing static and flying qubits like quantum memories, repeaters and interconnections between quantum computing units [2317].

Trapped-ions

Ion qubits operations can also be optimized with regards to the energetics of their gates or even for implementing a half-adder [2318].

Cold atoms

As we will see later in this book, significant progress was accomplished with cold atom FTQC computing since 2023.

In analog simulation mode, cold atoms could bring some potential quantum computational and energetic advantage. It is investigated by Pasqal [2319].

In FTQC gate-based mode, the technology could scale to as many as 10,000 atoms, thanks to improvements in atom controls. A 2024 Caltech experiment, published in *Nature* in 2025, controlled the position of 6,100 atoms [2320] and a team in China reached 11,000 trapped atoms in 2026 [2321]. FTQC experiments done by QuEra/Harvard in 2025 implemented all logical gate operations [1135]. However, the technology still relies on non-QND measurement, which leads experiments to use postselection, a non-scalable technique for FTQC. Many solutions are envisioned to enable QND measurement like using dual-species atoms.

Cold atoms computing benefits from a 1-to-many qubit connectivity enabled by atom shuttling. It powers efficient “non local” error correction codes with a high encoding rate, in other words a low physical-to-logical overhead. This has however a time and atom losses cost. Scaling cold atoms computing will also need efforts on light forming devices as well as on the development of low phase-noise high-power lasers. Qubit operation fidelities increased using 4 K cryogenics, which improved the quality of ultra-vacuum in the experiment chamber.

When all these questions are addressed, cold atom systems may sit in the low-power range of FTQC systems due to a good combination: large scale qubit control in a single system, 4K cryogenics, and modest energetic cost per qubit control. Undisclosed industry vendors plans are to build large scale FTQC system under the 1 MW power consumption threshold.

For more on this, look at [2322, 2323] which provide some clues on how to measure the energetic cost of cold atoms quantum systems.

Photon qubits

These are different beasts with regards to the thermodynamics of the whole value chain between single photon generation, entanglement preparation, computing (usually, using MBQC), spin-photon interfaces and single photon readout [2317].

Other news drive less enthusiasm like PsiQuantum’s 2024 preprint, published in *Nature* in 2025, and the related company communication, which help infer the power consumption of their 100 logical qubit system that uses an integrated silicon photonic circuit operating at 1.8 K, with thermally controlled phase shifters [2204]. Each cabinet holding PsiQuantum’s processors would need 360 W of cooling at a 2K temperature. Four such units will be built, one in Brisbane in Australia and three in Chicago (plus an additional prototype in Chicago). We assume here that the PsiQuantum datacenter site power needs will be mostly allocated to powering these FTQC systems. The only good news is that this power should be the same for future generation QPUs supporting thousands logical qubits. What drives these energetic costs? Mainly the inefficiency of the heralded silicon photon source and the need to cool the whole chip that also includes wave guides and

superconducting detectors, along with cryo-electronics for controlling it. Xanadu has similar gigantic plans, expecting to build quantum data center across multiple skyscrapers and requiring tens of thousands of racks [2205].

The picture looks better with Quandela and Sparrow Quantum, due to their deterministic photon sources. It lowers the hardware overhead to implement either a measurement-based FBQC model [2324–2326] or the SPOQC Quandela architecture [1906]. Quandela applied the MNR methodology to Boson Sampling and found an energetic advantage per sample showing up *before* any computational advantage, in a regime where the best classical algorithms are still faster [2327]. Optical computing energetics has been a research topic for Pascale Senellart’s C2N team and Alexia Auffèves’s Majulab team for a long time [216]. It’s a rare field of exploration of quantum energetics to improve the efficiency of quantum dots photon sources and spin-photons interactions [2250, 2328–2333].

The energetics of photonic computing is also investigated in Brazil [2334], and in the field of photonic-based quantum reservoir computing [2335, 2336]. We also have the H2020 OPTologic project aimed at creating light-induced and controlled topology for energy-efficient logic operations in quantum photonic computing systems. And the French startup Qluster that works on continuous variables room-temperature photonic computing, with a very low energy footprint, although with a relatively low TRL.

5.4.5 Actions

Scaling quantum computing is one of the most challenging scientific and technology endeavors ever launched by mankind on top of space exploration, nuclear fusion, DNA sequencing and genome based therapies creations. It should be undertaken with behaving responsibly as early as possible. One way is to embed in research and systems design an approach integrating the environmental footprint of quantum technologies. This footprint is of course energetic but also encompasses raw materials, manufacturing processes and product lifecycle estimations and optimizations [2337]. Addressing these questions is at once a scientific, a technological and a societal challenge.

Learning from AI

We can learn a couple lessons from what happened with artificial intelligence and deep learning. It became trendy starting in 2012 with a peak around 2020 when deep learning use cases became mainstream and embedded from smartphones to cloud datacenters. It was then discovered that AI had a significant energetic cost, both for training large deep learning models and to run them whether on end-user devices or on servers [2148]. The “frugal AI” topic then emerged. Solutions were proposed to reduce the energetic footprint of AI mainly with less

data-hungry machine learning models [2338], so-called data “quantization” (using 8-bit and even 1-bit numbers instead of 16-32-64 floating-point numbers) and with optimizing the power consumption of dedicated hardware like GPGPUs and embedded systems chips (in smartphones, connected objects and cars). Even in the LLM landscape, continuous progress is made starting at the academic and industry levels to reduce the cost of both training and inferences [2339–2346]. What if the environmental footprint of AI had been taken care of earlier?

Likewise, the environmental footprint of semiconductor production is questioned. It seems significant in energy and water consumption, and also on pollutants use [2347].

Quantum Energy Initiative

The same questions deserve to be asked about quantum technologies. Why not take care right now of their environmental footprint? One could argue that the first challenge is scientific before being environmental. Some are advocating to create high-fidelity qubits and useful fault-tolerant quantum computers first, and later address their related environmental aspects. Looking at how research labs and industry vendors were working until now on addressing the scalability challenges of quantum computers demonstrates that despite its relative technology immaturity and high scientific uncertainty, it is time to take environmental concerns into account right now. In a world of doubts on the role of science and technology, it is also a way to demonstrate that in emerging technologies, it is possible to implement responsible innovation practices from the start and not as afterthoughts and under pressure [2350].

This is the reasoning behind the creation of the Quantum Energy Initiative (QEI) in 2022 [2351].



The idea came out from a research team led by Alexia Auffèves (CNRS MajuLab Director, Singapore) with Robert Whitney (CNRS LPMMC in Grenoble), along with Janine Splettstoesser (Chalmers University) and Olivier Ezratty (this book’s author).

The QEI tries to answer several key questions related to quantum computing:

- Is there a **quantum energy advantage** as the processors scale up and how different is it from the quantum computational advantage?
- What is the fundamental **minimal energetic cost** of quantum computing?
- How to **avoid energetic dead-ends** on the road to large scale quantum computing? Can we create optimization tools and models for qubit technology, enabling technologies and software engineering?

The seed of Quantum Energy Initiative is described in a June 2022 PRX Quantum perspective [2352]. It lays the ground for a transversal initiative, connecting quantum thermodynamics, quantum information science, quantum physics and engineering. It makes the connection between classical and quantum thermodynamics, qubit architectures, qubit noise models, room temperature control electronics and cryo-electronics, quantum error correction codes, algorithms and compiler designs. It proposes a methodology to assess the energetic performance of quantum technologies, dubbed MNR [1734].

Given the diversity of research questions and domains, they require a broad inter-disciplinary community using numerous methodologies in numerous research projects. The Quantum Energy Initiative was created with the goal of bottom-up community-building, to unite such an inter-disciplinary community around the question of energy and energy-efficiency in all quantum technologies. A central goal of this community is to develop common languages for “energy”. For example, its aim is that quantum physicists and enabling-technology engineers can speak the same language and understand the relationships between energy at the quantum level (the level of individual qubits) and the macroscopic level (the level of cryogenics, laser and electronics). This requires encouraging exchange within this growing community, so that community members can understand each other’s research challenges. The Quantum Energy Initiative’s activities are also there to allow scientists and engineers to find potential collaborators with complementary interests and allow companies to recruit young people with expertise on energy in quantum technologies. A secondary goal of the Quantum Energy Initiative is to raise awareness of questions of energy-efficiency in quantum technologies among funding bodies, so that high-quality research projects on these questions have a fair chance of getting funded.



Figure 302: The group photo of QEI 2026 in Barcelona, Spain, in May 2026, which gathers 40 speakers and 112 participants.

This bottom-up community is centered around a website and newsletter, and lives entirely through the voluntary contribution of its members. As of September 2026, 34 academic and industry partners had joined and endorsed the QEI, including IBM, Microsoft, IQM, Riverlane, Alice&Bob, Quandela, Pasqal and Quobly, hybrid HPC-quantum services and others. It has a scientific



	atoms		electrons: controlled spin and microwave cavities			photons
						
	cold atoms	trapped ions	superconducting	silicon spin	NV centers	photons
cryogenics	7 to 10 kW for the 4 K cryostat cooling the science chamber, needed above a few 100s atoms [2323]	300 W to 6 kW, none for room-temperature traps, 4 to 15 K for cryogenic traps improving vacuum and heating rates [2295]	16 kW for one dilution fridge hosting ≈ 50 qubits, ≈ 105 kW for a Bluefors KIDE with 9 pulse tubes and 3 dilutions [2292, 2348]	12 kW, the 100 mK to 1.5 K qubit operating point relaxing the thermal budget [2185, 2349]	<1 kW, room temperature for single-qubit control, 1 to 4 K for two-qubit gates	3 kW for the 2 to 4 K stages cooling the single-photon sources and the SNSPD detectors [2205]
vacuum pumps	ultra-high vacuum, ≈ 100 W for the pump	ultra-high vacuum, chamber bake-out with heating strips over days to weeks, then stopped and stable during computation	vacuum pump	vacuum pump	none at room temperature, vacuum only when operated at 4 K	none for photonic circuits, cryostat vacuum only for sources and detectors
qubit gates control	1.8 kW for atom heaters, cooling and Rydberg lasers, SLM and AOD trap control, readout camera and electronics, roughly fixed up to about 1,000 atoms [2322, 2323]	<1.4 kW for ion heating, lasers, microwave generation and CMOS readout electronics, the RF trap drive alone costing about 3 W per ion [2295, 2318]	20 mW to 100 W per qubit depending on multiple parameters [1734, 2279]	20 mW to 100 W per qubit, cryo-CMOS controllers demonstrated at 3 K in the mW per qubit range [2279, 2349]	N/A, no register large enough to publish a per-qubit control budget	300 W for photon sources, detectors and phase-shifter control, about 7.5 W per qubit for detection against 1 mW per qubit for generation [1906, 2334]
classical computing	300 W in NISQ	300 W in NISQ	<1 kW, plus real-time syndrome decoding at 8 mW for an ASIC decoder against a few W for its FPGA equivalent [2266]	<1 kW, plus real-time syndrome decoding [2266]	<1 kW	700 W, higher because feed-forward and fusion outcomes must be processed at the photon clock rate
# qubits used	100 (Pasqal Fresnel) and 256 (QuEra Aquila) (1), 300 to 1,000 in next-generation cryogenic systems (2)	20 to 24 (AQT), up to 98 (Quantinuum Helios) (5)	53 (Google Sycamore) to 3×156 (IBM Quantum System Two with Heron r2) (3)	12 (Intel Tunnel Falls)	<8	20 modes (Quandela, QuiX) (4)
total	3 kW (1) to 20 kW (2)	2 to 34 kW (5)	25 to 200 kW (3)	21 kW	N/A	4 kW (4)

Table 12: Rough estimations of current NISQ quantum computers total power and its decomposition. The figures mix measured values, vendor statements, modeled values and author inferences from assumed configurations. Typical configurations for Pasqal and QuEra (1), neutral atoms with 4 K cryogenics (2), Google 53-qubit Sycamore and guesstimate for IBM Quantum System Two with its Bluefors KIDE cryostat (3), Quandela and QuiX (4), AQT and Quantinuum (5), rough estimates for the others. You must also consider the difference in speed of operations between these QPUs and their qubit connectivity, which affects the circuit run times and thus their energy efficiency. Per-qubit control figures and total system power do not multiply out, since control electronics is largely a fixed cost at these qubit counts for atoms, ions and photons, and only becomes proportional to the qubit count above a few hundred qubits. © Olivier Ezratty, November 2023-September 2026.

board covering 5 continents and has launched its own video seminar series.

The QEI also organized its first one-week workshops in Singapore in November 2023, Grenoble in January 2025, Barcelona in May 2026 (Figure 302), with the 2027 edition planned in Lisbon in September 2027.

The QEI garnered visibility in various quantum related events and even in Nature [2353]. It also led to the creation of the IEEE Working Group for creating a Standard for Quantum Computing Energy Efficiency [2354]. It may lead to the creation of a Quantum Green ranking of QPUs, a bit like the Green500 in HPC [2355].

The QEI is not limited to quantum computing and is intended to expand to all quantum technologies, namely quantum telecommunications [2356–2358] and quantum sensing.


A separated but connected topic is the use cases of quantum computing which may address climate change one way or the other. This is the “supply side” of the energetic question [2359].

Other initiatives

The QEI is not the only organization caring about the energetics of quantum technologies. OQI and GESDA are also involved as well as the World Economic Forum. Quantum technologies responsible innovation is described in a paper by Dakota Root, then a post-doc at UGA in Grenoble and now in the USA [2360]. Her recommendations are about education, creating a supply chain sustainability mandate, funding research on minimizing carbon footprint of quantum technologies, and develop standards for evaluating it. Key questions are:

is there a quantum computing energetic advantage? For solving a problem, is this the best quantum computer for the task? Is the energy consumption for this task

reasonable? Are the quantum computer's parameters tuned to run in a frugal manner?



5.4.6 Bibliography and notes

These are the bibliographical references and notes for the **Energetics** subsection.

- [216] [Experimental analysis of energy transfers between a quantum emitter and light fields](#), by Ilse Maillette de Buy Wenniger, Alexia Auffèves, Pascale Senellart, et al., arXiv, June 2023 (18 pages).
- [236] [Phase transition in Random Circuit Sampling](#), by A. Morvan et al., arXiv, April 2023 (43 pages).
- [475] [Quantum speed limits: from Heisenberg's uncertainty principle to optimal quantum control](#), by Sebastian Deffner and Steve Campbell, Journal of Physics A: Mathematical and Theoretical, October 2017 (50 pages).
- [754] [Powering Quantum Computation with Quantum Batteries](#), by Yaniv Kurman, Physical Review X, January 2026 (17 pages).
- [1135] [A fault-tolerant neutral-atom architecture for universal quantum computation](#), by Dolev Bluvstein, Mikhail D. Lukin, et al., Nature, November 2025 (30 pages).
- [1734] [Optimizing resource efficiencies for scalable full-stack quantum computers](#), by Marco Fellous-Asiani, Jing Hao Chai, Yvain Thonnart, Hui Khoo Ng, Robert S. Whitney, and Alexia Auffèves, arXiv, October 2023 (41 pages).
- [1820] [A Scalable FPGA Architecture for Real-Time Decoding of Quantum LDPC Codes Using GARI](#), by Daniel Báscones, Arshpreet Singh Maan, Valentin Savin, and Francisco Garcia-Herrero, arXiv, May 2026 (11 pages).
- [1856] [Logical quantum processor based on reconfigurable atom arrays](#), by Dolev Bluvstein, Mikhail D. Lukin, et al., Nature, December 2023 (28 pages).
- [1906] [A Spin-Optical Quantum Computing Architecture](#), by Grégoire de Glinasty, Paul Hilaire, Pierre-Emmanuel Emeriau, Stephen C. Wein, Alexia Salavrakos, and Shane Mansfield, arXiv, July 2024 (29 pages).
- [1956] [Optimization by Decoded Quantum Interferometry](#), by Stephen P. Jordan, Ryan Babbush, et al., Nature, October 2025 (7 pages).
- [2137] [Measuring the Emissions and Energy Footprint of the ICT Sector: Implications for Climate Action](#), by World Bank, World Bank, March 2024 (54 pages).
- [2138] [Electricity Consumption](#), by IEA, IEA, September 2024.
- [2139] [Data centres and data transmission networks](#), by International Energy Agency, IEA tracking report, January 2024.
- [2140] [Executive summary – Electricity 2025 – Analysis - IEA](#), by IEA, IEA, February 2025.
- [2141] [Electricity 2026: Analysis and forecast to 2030](#), by International Energy Agency, IEA report, Paris, February 2026.
- [2142] [Key Questions on Energy and AI](#), by International Energy Agency, IEA report, Paris, April 2026.
- [2143] [ICT sector electricity consumption and greenhouse gas emissions - 2020 outcome](#), by Jens Malmodyn, Nina Lovehagen, Pernilla Bergmark, and Dag Lunden, Telecommunications Policy, April 2024 (23 pages) 📄.
- [2144] [Data on AI Data Centers](#), Epoch AI, August 2026.
- [2145] [ICT Sector Greenhouse Gas Emissions – Issues and Trends](#), by Peter Garraghan, John Hutchinson, and Adrian Friday, arXiv, October 2024 (13 pages).
- [2146] [Green ICT](#), by Informatics Europe, 2026.
- [2147] [The Need for a Green ICT Reference Framework](#), by Marco Aiello, Mina Alipour, Antonio Brogi, Rafael Capilla, Lidia Fuentes, Ilche Georgievski, Gabriele Gianini, Mahyar Tourchi Moghaddam, Monica Vitali, and Sebastian Werner, arXiv, April 2026.
- [2148] [Compute and Energy Consumption Trends in Deep Learning Inference](#), by Radosvet Desislavov, Fernando Martínez-Plumed, and José Hernández-Orallo, arXiv, March 2023 (26 pages).
- [2149] [Position: Tensor Networks are a Valuable Asset for Green AI](#), by Eva Memmel, Clara Menzen, Jetze Schuurmans, Frederiek Wesel, and Kim Batselier, arXiv, May 2024 (14 pages).
- [2150] [Rethinking Concerns About AI's Energy Use](#), by Daniel Castro, Center for Data Innovation, January 2024 (22 pages).
- [2151] [Scaling Laws for Energy Efficiency of Local LLMs](#), by Ander Alvarez, Alessandro Genuardi, Nilotpal Sinha, Antonio Tiene, Mikail Okay, Bakbergen Ryskulov, David Montero, Samuel Mugel, and Román Orús, arXiv, December 2025 (9 pages).
- [2152] [The growing energy footprint of artificial intelligence](#), by Alex de Vries, Cell, 2023.
- [2153] [The Economics of Digital Intelligence Capital: Endogenous Depreciation and the Structural Jevons Paradox](#), by Yukun Zhang and Tianyang Zhang, arXiv, January 2026 (46 pages).
- [2154] [From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate](#), by Alexandra Sasha Luccioni, Emma Strubell, and Kate Crawford, FAccT '25: Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency, June 2025 (13 pages).
- [2155] [Trends in AI Supercomputers](#), by Konstantin F. Pilz, James Sanders, Robi Rahman, and Lennart Heim, arXiv, April 2025 (48 pages).
- [2156] [An Initial Exploration of Contrastive Prompt Tuning to Generate Energy-Efficient Code](#), by Sophie Weidmann and Fernando Castor, arXiv, March 2026 (8 pages).
- [2157] [Power Hungry Processing: Watts Driving the Cost of AI Deployment?](#), by Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell, arXiv, October 2024 (21 pages).
- [2158] [How Much Energy Do LLMs Consume? Unveiling the Power Behind AI](#), by Sourabh Mehta, Association of Data Scientists, July 2024.
- [2159] [Evaluating the Energy Consumption of Machine Learning: Systematic Literature Review and Experiments](#), by Charlotte Rodriguez, Giovanni Neglia, et al., arXiv, August 2024 (52 pages).
- [2160] [10 Largest AI Data Centers in the World](#), by Tamzid, Brightlio - Technology Illuminated, July 2026.
- [2161] [Amazon CEO Andy Jassy talks 6 truths surrounding the rise of AI](#), by Amazon Staff, Amazon News, April 2026.
- [2162] [Amazon plans to invest \\$15 billion in Northern Indiana to build new data center campuses and advance AI innovation](#), by Amazon Staff, Amazon News, November 2025.
- [2163] [ICT energy evolution: telecom, data centers, and AI](#), by Daniel Paska, Nina Lovehagen, Ove Persson, and Jens Malmodyn, Ericsson white paper, April 2025 (12 pages).

- [2164] [TOP500 List, June 2026, 67th edition](#), by TOP500.org, TOP500 project, announced at ISC 2026, Hamburg, June 2026.
- [2165] [Life cycle analysis for emissions of scientific computing centres](#), by Mattias Wadenstein and Wim Vanderbauwhede, The European Physical Journal C, August 2025 (14 pages).
- [2166] [ACT: designing sustainable computer systems with an architectural carbon modeling tool](#), by Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu, Proceedings of the 49th Annual International Symposium on Computer Architecture (ISCA), pages 784-799, June 2022 (16 pages) .
- [2167] [Technology Carbon Standard: data centre and server hardware lifecycle example](#), by Technology Carbon Standard, Technology Carbon Standard, open methodology, January 2025.
- [2168] [Global server market shipments and revenue](#), by TrendForce and IDC, TrendForce and IDC quarterly server trackers, compiled, February 2024.
- [2169] [Invisible Emissions: a forecast of tech supply chain emissions and electricity consumption by 2030](#), by Greenpeace East Asia, Greenpeace East Asia report, April 2023.
- [2170] [Teachers and Local Businesses Win as Meta Expands Louisiana Data Center](#), by Meta, Meta Newsroom, July 2026.
- [2171] [Elon Musk says xAI will increase data center capacity 7x by 2027 — targeting 10 gigawatts of compute, up to \\$500 billion in revenue by the end of next year](#), by Anton Shilov, Tom's Hardware, August 2026.
- [2172] [This Is How the AI Bubble Will Pop](#), by Derek Thompson, Derek Thompson website, October 2025.
- [2173] [Mistral AI wants to build 1 gigawatt of European compute by 2030 — and lock in customers now](#), by Michael Nuñez, VentureBeat, August 2026.
- [2174] [Nebius to construct 310 MW AI factory in Finland](#), by Nebius, Nebius website, March 2026.
- [2175] [China's data center capacity set to top 60 GW by 2030, driving a doubling of power demand](#), by Simeng Deng, Fei Chen, Minh Khoi Le, and Kartik Selvaraju, Rystad Energy, April 2026.
- [2176] [NVIDIA Computex 2024 Keynote NVIDIA Rubin GPU Vera CPU and 1.6T Networking](#), by Patrick Kennedy, June 2024.
- [2177] [Is Quantum Computing An Unlikely Answer To AI's Looming Energy Crisis?](#), by Tor Constantino, Forbes, October 2024.
- [2178] [Bitcoin's Energy consumption: a quantitative approach on a subjective question](#), by Rachek Rybarczyk, Galaxy Digital Mining, May 2021 (13 pages).
- [2179] [Fact sheet: Climate and Energy Implications of Crypto-Assets in the United States](#), by White House OSTP, White House, September 2022.
- [2180] [Ethereum energy consumption](#), by Ethereum, Ethereum, September 2022.
- [2181] [Bitcoin Energy Consumption Index](#), by Digiconomist, Digiconomist, July 2024.
- [2182] [IEA Study Sees AI, Cryptocurrency Doubling Data Center Energy Consumption by 2026](#), by Sean Buckley, Data Center Frontier, March 2024.
- [2183] [Ranked: Top 50 Data Center Markets by Power Consumption](#), by Julie R. Peasley, Visual Capitalist, January 2024.
- [2184] [Key Questions on Energy and AI: Executive Summary](#), by Thomas Spencer, Siddharth Singh, and Laura Cozzi, April 2026 (6 pages).
- [2185] [Energy use in quantum data centers: Scaling the impact of computer architecture, qubit performance, size, and thermal parameters](#), by Michael James Martin, Ray Grout, et al., IEEE Transactions on Sustainable Computing, October 2022 (12 pages)  (preprint on [osti.gov](#)).
- [2186] [Energy Efficiency trends in HPC: what high-energy and astrophysicists need to know](#), by Estela Suarez, Jorge Amaya, Martin Frank, Oliver Freyermuth, Maria Girona, Bartosz Kostrzewa, and Susanne Pfalzner, arXiv, March 2025 (45 pages).
- [2187] [A Bidirectional Gated Recurrent Unit Model for PUE Prediction in Data Centers](#), by Dhivya Dharshini Kannan, Anupam Trivedi, and Dipti Srinivasan, arXiv, December 2025 (8 pages).
- [2188] [Implications of Historical Trends in the Electrical Efficiency of Computing](#), by Jonathan Koomey et al., IEEE Annals of the History of Computing, February 2011 .
- [2189] [Moore's Law Might Be Slowing Down, But Not Energy Efficiency](#), by Jonathan Koomey and Samuel Naffziger, IEEE Spectrum, March 2015.
- [2190] [Evolution of computing energy efficiency: Koomey's law revisited](#), by Alberto Prieto, Beatriz Prieto, Juan José Escobar, and Thomas Lampert, Cluster Computing, October 2024 (22 pages).
- [2191] [Top 500 - The list - June 2025](#), by Top500.org, June 2025.
- [2192] [Ranking programming languages by energy efficiency](#), by Rui Pereira, Science, January 2021 (30 pages).
- [2193] [Photonics for Sustainable Computing](#), by Farbin Fayza, Satyavolu Papa Rao, Darius Bunandar, Udit Gupta, and Ajay Joshi, arXiv, January 2024 (4 pages).
- [2194] [What limits the simulation of quantum computers?](#), by Yiqing Zhou et al., PRX, November 2020 (14 pages).
- [2195] [A density-matrix renormalization group algorithm for simulating quantum circuits with a finite fidelity](#), by Thomas Ayrat, Thibaud Louvet, Yiqing Zhou, Cyprien Lambert, E. Miles Stoudenmire, and Xavier Waintal, arXiv, August 2022 (25 pages).
- [2196] [Achieving Energetic Superiority Through System-Level Quantum Circuit Simulation](#), by Rong Fu, Wanli Ouyang, et al., arXiv, June 2024 (11 pages).
- [2197] [Leapfrogging Sycamore: Harnessing 1432 GPUs for 7× Faster Quantum Random Circuit Sampling](#), by Xian-He Zhao, Ming-Cheng Chen, et al., arXiv, June 2024 (8 pages).
- [2198] [Energy-Efficient Quantum Computing](#), by Joni Ikonen, Juha Salmilehto, and Mikko Möttönen, arXiv, September 2016 (12 pages).
- [2199] [Energy-error tradeoff in encoding quantum error correction](#), by Josey Stevens and Sebastian Deffner, arXiv, May 2026 (8 pages).
- [2200] [Top 500 - The list - June 2024](#), by Top500.org, June 2024.
- [2201] [Roll call purchase recommendations for Illinois Quantum and Microelectronics Park](#), by IQMP Board of Trustees, November 2024 The document mentions that “*IQMP, LLC will then execute a \$15,000,000 purchase order with ComEd for a new onsite substation with multiple large transformers and capacity to supply 350MW of power at the IQMP*”. These include 300MW to cover the need for three PsiQuantum systems with 100MW power draw each, and 50 MW for the rest of the IQMP park which hosts the National Quantum Facility (NQF). NCF will provide cryogenic services with up to 20 traditional cryostats for the park participants like IBM, Pasqal, Infleqtion and Diraq.
- [2202] [Powering the Future: How ComEd's Partnership Sparked the Illinois Quantum Campus](#), by Excelon, Excelon website, August 2025 This press release from the electricity provider of the ICMP site in Chicago that will host three PsiQuantum computers mentions that: “*The demands of powering a utility-scale quantum campus are immense, estimated at approximately 350 megawatts, enough to support not only PsiQuantum but also shared facilities like a state-of-the-art cryogenic plant and, eventually, other cutting-edge companies drawn to the site [...]. Meeting these needs required out-of-the-box thinking. When PsiQuantum decided to accelerate the development of a prototype quantum computer at the IQMP, ComEd responded by repurposing retired transmission infrastructure, enabling rapid*

delivery of up to 100 megawatts". We infer from this that a single PsiQuantum system needs 100 MW and that three of them will be installed in Chicago in their quantum data center.

- [2203] We can decompose the 100 MW needed to power a PsiQuantum computer with the following reasoning. It will assemble a couple hundred cabinets containing hundreds of chips. Each cabinet will need 360 W of cooling power at about 2 K according to PsiQuantum. Cooling will be provided by a mix of liquid helium (for 4.5 K) and cold-boxes (from 4.5 K to 2 K). The specific power need of the whole cryogenic cycle will be similar to the one from the CERN LHC and the Facility for Rare Isotope Beams of the University of Michigan, between 500 and 800 W/W. This alone creates a cryogenic power need sitting between 40 and 60 MW. Then, one must add all room-temperature control electronics for the photonic chips, laser sources, and quantum error correction classical computing resources.
- [2204] [A manufacturable platform for photonic quantum computing](#), by PsiQuantum team, Nature, February 2025 (9 pages).
- [2205] [Scaling and networking a modular photonic quantum computer](#), by H. Aghaee Rad, Y. Zhang, et al., Nature, January 2025.
- [2206] [Decode Quantum with Christian Weedbrook from Xanadu](#), by Olivier Ezratty, Opinions Libres, 2026.
- [2207] [Xanadu 2026 Analyst Day](#), by Xanadu, Xanadu investor relations, March 2026 slide 26.
- [2208] [Beyond Quantum Advantage: The Future of Computing](#), by Jay Gambetta, SCA/HPCAsia 2026, January 2026 the presentation includes a slide with the title "Modular and efficient cryogenics" dedicated to resource estimates for their upcoming Starling and Bluejay FTQC systems. It says "Power* (2-gross): 1.5 M for Starling, 6.3 MW for Blue Jay", * precising that this includes cryogenics, control electronics, and classical runtime. Starling and Blue Jay cryogenics will be different than with IBM System Two which uses a Bluefors KIDE. They will use centralized liquid helium production at about 4 K that will be complemented by dilutions in the cryostats to reach the 10-15 mK required for qubit operations. The presentation video was not made public but the conference was public.
- [2209] Power information provided by the company to the author, with the authorization to embed it in the chart, in September 2026. The circuit depth doesn't influence the power but only computing time and the total energy consumed to run a circuit.
- [2210] [Could Quantum Fix AI's Energy Problem?](#), by Diraq, The Spin on Quantum by Diraq (Substack), September 2026.
- [2211] [The case for silicon - How Diraq is making computers quantum](#), by Diraq, arXiv, September 2026 (25 pages).
- [2212] [Alice & Bob roadmap \(Graphene\)](#), by Alice & Bob, Alice & Bob website as of July 28th, 2026, July 2026.
- [2213] [Quantinuum breakout year: redefining global quantum leadership](#), by Mark Jackson, Q2B Santa Clara 2025, December 2025 (13 mn video).
- [2214] [C12 roadmap](#), by C12, C12 website, April 2026 Their 792 logical qubits are to be supported by 100,000 physical qubits, by 2033. This roadmap was announced in April 2026.
- [2215] [What Building the Next Computing Paradigm Actually Requires](#), by QuEra, QuEra web site, January 2026.
- [2216] [Energy Cost of Quantum Circuit Optimisation: Predicting That Optimising Shor's Algorithm Circuit Uses 1 GWh](#), by Alexandru Paler and Robert Basmadjian, ACM Transactions on Quantum Computing, March 2022 📄.
- [2217] [Le parc des réacteurs nucléaires français](#), by IRSN, Quantum Science and Technology, February 2022.
- [2218] [Small Nuclear Power Reactors](#), by World Nuclear Association, Quantum Science and Technology, February 2024.
- [2219] [From quantum speed limits to energy-efficient quantum gates](#), by Maxwell Aifer and Sebastian Deffner, arXiv, February 2022 (19 pages).
- [2220] [Fundamental speed limits on entanglement dynamics of bipartite quantum systems](#), by Vivek Pandey et al., Physical Review A, July 2023 📄 (preprint on [arXiv](#)).
- [2221] [Exact Quantum Speed Limits](#), by Arun K. Pati, Brij Mohan, Sahil, and Samuel L. Braunstein, arXiv, August 2023 (9 pages).
- [2222] [A Unifying Quantum Speed Limit For Time-Independent Hamiltonian Evolution](#), by H. F. Chau and Wenxin Zeng, arXiv, May 2024 (22 pages).
- [2223] [Quantum speed limit for complex dynamics](#), by Mao Zhang, Huai-Ming Yu, and Jing Liu, npj Quantum Information, October 2023 📄 (preprint on [dx.doi.org](#)).
- [2224] [Chaotic fluctuations in a universal set of transmon qubit gates](#), by Daniel Basilewitsch, Simon-Dominik Börner, Christoph Berke, Alexander Altland, Simon Trebst, and Christiane P. Koch, arXiv, January 2024 (5 pages).
- [2225] [Quantum Speed Limits for Implementation of Unitary Transformations](#), by Abolfazl Farmanian and Vahid Karimpour, arXiv, June 2024 (13 pages).
- [2226] [Roadmap on Quantum Thermodynamics](#), by Steve Campbell, Irene D'Amico, Mario A. Ciampini, Janet Anders, Natalia Ares, Alexia Auffèves, Michele Campisi, Sebastian Deffner, Cyril Elouard, Mauro Paternostro, Jukka Pekola, Janine Splettstoesser, Nicole Yunger Halpern, et al., arXiv, January 2026 (65 pages).
- [2227] [Universal Foundations of Thermodynamics: Entropy and Energy Beyond Equilibrium and Without Extensivity](#), by Gian Paolo Beretta, arXiv, March 2026 (84 pages).
- [2228] [A short story of quantum and information thermodynamics](#), by Alexia Auffèves, arXiv, March 2021 (14 pages).
- [2229] [Third law of thermodynamics and the scaling of quantum computers](#), by Lorenzo Buffoni et al., Physical Review Letters, October 2022 📄 (preprint on [arXiv](#)).
- [2230] [Nonequilibrium boundary-driven quantum systems: Models, methods, and properties](#), by Gabriel T. Landi et al., Review of Modern Physics, December 2022 (59 pages) 📄 (preprint on [arXiv](#)).
- [2231] [Quantum Thermodynamics](#), by Patrick P. Potts, arXiv, June 2024 (87 pages).
- [2232] [The Jaynes Cummings model as an autonomous Maxwell demon](#), by Yashovardhan Jha, Dragi Karveski, and Cyril Elouard, arXiv, February 2025 (14 pages).
- [2233] [Weak continuous measurements require more work than strong ones](#), by Lorena Ballesteros Ferraz and Cyril Elouard, arXiv, March 2025 (21 pages).
- [2234] [Thermodynamics of Hamiltonian anyons with applications to quantum heat engines](#), by Joe Dunlop, Álvaro Tejero, Michalis Skotiniotis, and Daniel Manzano, arXiv, March 2025 (21 pages).
- [2235] [Extracting energy via bosonic Gaussian operations](#), by Frank Ernesto Quintela Rodriguez, Francesco Anna Mele, Salvatore Francesco Emanuele Oliviero, Vittorio Giovannetti, Ludovico Lami, and Vasco Cavina, arXiv, March 2025 (21 pages).
- [2236] [Learning to erase quantum states: thermodynamic implications of quantum learning theory](#), by Haimeng Zhao, Yuzhen Zhang, and John Preskill, arXiv, April 2025 (5 pages).
- [2237] [The Quantum Toll Framework: A Thermodynamic Model of Collapse and Coherence](#), by L. S. Montejo, arXiv, May 2025 (33 pages).
- [2238] [Revisiting thermodynamics in computation and information theory](#), by Pritam Chattopadhyay and Goutam Paul, arXiv, August 2024 (40 pages).
- [2239] [Landauer Principle and Thermodynamics of Computation](#), by Pritam Chattopadhyay, Avijit Misra, Tanmoy Pandit, and Goutam Paul, arXiv, August 2025 (41 pages).

- [2240] [Experimental Extraction of Coherent Ergotropy and Its Energetic Cost in a Superconducting Qubit](#), by Li Li, Silu Zhao, Kai Xu, Heng Fan, Dongning Zheng, and Zhongcheng Xiang, arXiv, June 2025 (7 pages).
- [2241] [Amplifying microwave pulses with a single qubit engine fueled by quantum measurements](#), by Rémy Dassonneville, Cyril Elouard, Romain Cazali, Réouven Assouly, Audrey Bienfait, Alexia Auffèves, and Benjamin Huard, Physical Review Research, March 2026 (14 pages).
- [2242] [Revealing the fuel of a quantum continuous measurement-based refrigerator](#), by Cyril Elouard, Sreenath K. Manikandan, Andrew N. Jordan, and Geraldine Haack, arXiv, February 2025 (18 pages).
- [2243] [Hamiltonian quantum gates – energetic advantage from entangleability](#), by Josey Stevens and Sebastian Deffner, arXiv, July 2025 (7 pages).
- [2244] [Quantifying entanglement in quantum thermodynamics via separability constraints](#), by Joan Alba, Laura Ares, Jan Sperling, and Julien Pinske, arXiv, March 2026 (14 pages).
- [2245] [Autonomous quantum heat engine](#), by Tuomas Uusnäkki, Miika Rasola, Vasilii Vadimov, Priyank Singh, Ahmad Darwish, and Mikko Möttönen, arXiv, March 2026 (16 pages).
- [2246] [Illusory cracks in the second law of thermodynamics in quantum nanoelectronics](#), by Robert S. Whitney, arXiv, April 2023 (75 pages).
- [2247] [Bosonic statistics enhance Maxwell’s demon in photonic experiment](#), by Malaquias Correa Anguita, Sara Marzban, William F. Braasch Jr., Twesh Upadhyaya, Gabriel Landi, Nicole Yunger Halpern, and Jelmer J. Renema, arXiv, February 2026 (5 pages).
- [2248] [Thermodynamic Proof of Quantumness without Structure](#), by Florian Meier and Hayata Yamasaki, arXiv, September 2026 (16 pages).
- [2249] [No-Go Theorem for Quantum Heat Engines Powered Purely by Quantum Measurements in the Steady Regime](#), by Kenta Koshihara and Kazuya Yuasa, arXiv, April 2026 (10 pages).
- [2250] [Energetics of non-Gaussianity in single mode cavities](#), by Sahil Sardar Jafar, Kian Hwee Lim, Kiarn T. Laverick, Samyak P. Prasad, Maria Maffei, and Alexia Auffèves, arXiv, June 2026 (11 pages).
- [2251] [Engineered Dissipation for Quantum Information Science](#), by Patrick M. Harrington, Erich Mueller, and Kater Murch, arXiv, May 2022 (32 pages).
- [2252] [Energy dynamics, heat production and heat-work conversion with qubits: towards the development of quantum machines](#), by Liliana Arrachea, arXiv, November 2022 (74 pages).
- [2253] [From basic science to technological development: the case for two avenues](#), by Matteo Carlesso and Mauro Paternostro, arXiv, May 2023 (21 pages).
- [2254] [Logical reversibility of computation](#), by Charles Bennett et al., IBM Journal of Research and Development, 1973 (8 pages).
- [2255] [Synthesis of Reversible Logic Circuits](#), by Vivek Shende, John P. Hayes, et al., 2002 (30 pages).
- [2256] [Temperature as Joules per Bit](#), by Charles Alexandre Bédard, Sophie Berthelette, Xavier Coiteux-Roy, and Stefan Wolf, arXiv, February 2025 (15 pages).
- [2257] [Pathfinding Thermodynamically Reversible Quantum Computation](#), by Karpur Shukla et al., January 2020 (28 slides).
- [2258] [Asynchronous Ballistic Reversible Computing using Superconducting Elements](#), by Michael P. Frank et al., Sandia National Laboratories, April 2020 (27 slides).
- [2259] [Quantum Energy Teleportation in Spin Chain Systems](#), by Masahiro Hotta, Journal of the Physical Society of Japan, February 2009 (22 pages)  (preprint on [arXiv](#)).
- [2260] [A review of applications of Quantum Energy Teleportation: from experimental tests to thermodynamics and space-time engineering](#), by Boris Ragula and Eduardo Martín-Martínez, arXiv, May 2025 (28 pages).
- [2261] [Scalable Repeater Architecture for Long-Range Quantum Energy Teleportation in Gapped Systems](#), by M. Y. Abd-Rabbou, Irfan Siddique, Saeed Haddadi, and Cong-Feng Qiao, arXiv, January 2026.
- [2262] [Quantum Computation with Quantum Batteries](#), by Yaniv Kurman, Kieran Hymas, Arkady Fedorov, William J. Munro, and James Quach, arXiv, April 2025 (15 pages).
- [2263] [IBM built the biggest, coolest quantum computer. Now comes the hard part](#), by Alex Pasternack, Fast Company, January 2024.
- [2264] [Thermal Capacity Mapping of Cryogenic Platforms for Quantum Computers](#), by Scott A. Manifold, George B. Long, and Jonathan J. Burnett, arXiv, March 2025 (12 pages).
- [2265] [Limitations in quantum computing from resource constraints](#), by Marco Fellous-Asiani et al., PRX Quantum, November 2021 (8 pages).
- [2266] [A real-time, scalable, fast and resource-efficient decoder for a quantum computer](#), by Ben Barber, Abbas B. Ziad, et al., Nature Electronics, January 2025  (preprint on [arXiv](#)).
- [2267] [EQISA: Energy-efficient Quantum Instruction Set Architecture using Sparse Dictionary Learning](#), by Sibashish Mishra, Aritra Sarkar, and Sebastian Feld, arXiv, March 2026 (24 pages).
- [2268] [Distributing Quantum Computations, Shot-wise](#), by Giuseppe Bisicchia, Giuseppe Clemente, Jose Garcia-Alonso, Juan Manuel Murillo Rodríguez, Massimo D’Elia, and Antonio Brogi, arXiv, November 2024 (22 pages).
- [2269] [The resource cost of large scale quantum computing](#), by Marco Fellous-Asiani, arXiv, November 2021 (215 pages).
- [2270] [Estimating the Energy Requirements to Operate a Cryptanalytically Relevant Quantum Computer](#), by Edward Parker and Michael J. D. Vermeer, arXiv, April 2023 (17 pages).
- [2271] [Consider the energy consumption of your quantum circuits](#), by Coral Calero, Macario Polo, Elena Desdentado, and Ma^a Ángeles Moraga, Nature Reviews Physics, June 2025 .
- [2272] [Exploring the trade-off between computational power and energy efficiency: An analysis of the evolution of quantum computing and its relation to classical computing](#), by Elena Desdentado, Coral Calero, Ma Ángeles Moraga, Manuel Serano, and Félix García, The Journal of Systems & Software, July 2024 (14 pages).
- [2273] [Optimizing resource allocation for accuracy in noisy variational quantum algorithms](#), by Harshit Verma, Thomas Ayrat, Alexia Auffèves, and Robert Whitney, arXiv, June 2026 (18 pages).
- [2274] [Quantum optimization algorithms: Energetic implications](#), by Rolando P. Hong Enriquez et al., Concurrency and Computation: Practice and Experience, March 2024 (21 pages)  (preprint on [hdl.handle.net](#)).
- [2275] [Uncertain quantum computing futures and potential energy and physical resource impacts at scale](#), by David L. McCollum, Naushita Sharma, Melrose Pan, Weijun Yao, Soumendra Bhanja, Sarah Chehade, and Andrea Delgado, Renewable and Sustainable Energy Transition, March 2026 (17 pages).
- [2276] [Helium resource global supply and demand: Geopolitical supply risk analysis](#), by Ankesh Siddhantakar et al., Resources, Conservation and Recycling, June 2023 .
- [2277] [Some Facts About Helium](#), by Walter Meissner Institute, Walter Meissner Institute web site, retrieved on July 29th, 2026, June 2023.
- [2278] [Optimizing resource efficiencies for scalable full-stack quantum computers](#), by Marco Fellous-Asiani, Jing Hao Chai, Yvain Thonnart, Hui Khoo Ng, Robert S. Whitney, and Alexia Auffèves, PRX Quantum, October 2023.
- [2279] [19.1 A Scalable Cryo-CMOS 2-to-20GHz Digitally Intensive Controller for 4×32 Frequency Multiplexed Spin Qubits/Transmons in 22nm FinFET Technology for Quan-](#)

- tum Computers, by Bishnu Patra, February 2020  (preprint on resolver.tudelft.nl).
- [2280] [Is quantum computing green? An estimate for an energy-efficiency quantum advantage](#), by Daniel Jaschke and Simone Montangero, Quantum Science and Technology, January 2023.
 - [2281] [The State of the Art in Integer Factoring and Breaking Public-Key Cryptography](#), by Fabrice Boudot et al., June 2022 (9 pages)  (preprint on [HAL](https://arxiv.org/abs/2206.00000)).
 - [2282] [The impact of hardware specifications on reaching quantum advantage in the fault tolerant regime](#), by Mark Webber, September 2021 (16 pages).
 - [2283] [Blueprint for a microwave trapped ion quantum computer](#), by Bjoern Lekitsch, Sebastian Weidt, Austin G. Fowler, Klaus Mølmer, Simon J. Devitt, Christof Wunderlich, and Winfried K. Hensinger, Science, February 2017 (11 pages).
 - [2284] [Could energy efficiency be quantum computers' greatest strength yet?](#), by Marco Fellous-Asiani, The Conversation, October 2022.
 - [2285] [In Search of Quantum Advantage: Estimating the Number of Shots in Quantum Kernel Methods](#), by Artur Mirowski, Marco Fellous Asiani, Jakub Mielczarek, Bertrand Le Saux, and Jakub Nalepa, arXiv, July 2024 (13 pages).
 - [2286] [Quantum Blockchain Miners Provide Massive Energy Savings](#), by Joseph Kearney and Carlos A Perez-Delgado, arXiv, June 2023 (6 pages).
 - [2287] [Quantum Miners: Revolutionizing Energy Efficiency in Blockchain](#), by Uvin Vindula, June 2023.
 - [2288] [The lower energy consumption in cryptocurrency mining processes by SHA-256 Quantum circuit design used in hybrid computing domains](#), by Ahmet Orun and Fatih Kurugollu, arXiv, December 2023 (10 pages).
 - [2289] [Quantum Energy Advantage](#), by BTQ, March 2024.
 - [2290] [Energy-Consumption Advantage of Quantum Computation](#), by Florian Meier and Hayata Yamasaki, PRX Energy, May 2025 (51 pages).
 - [2291] [Potential Energy Advantage of Quantum Economy](#), by Junyu Liu, Hansheng Jiang, and Zuo-Jun Max Shen, arXiv, August 2023 (23 pages).
 - [2292] [Supplementary information for “Quantum supremacy using a programmable superconducting processor”](#), by Frank Arute, Hartmut Neven, John M. Martinis, et al., arXiv, December 2019 (67 pages).
 - [2293] [Quantum Technology Monitor](#), by Henning Soller et al., McKinsey, April 2026 (103 slides).
 - [2294] [Energy efficiency of quantum computers](#), by Miquel Carrasco-Codina, Pau Escofet, Paul Hilaire, Ariane Soret, Sam Nerenberg, Victor Champain, Gerard Milburn, Klara Theophilo, Sophie H. Li, Irais Bautista, Andrés Gómez, Jose Miralles, Sergi Abadal, Carmen G. Almudéver, Eduard Alarcón, and Raja Yehia, arXiv, May 2026 (66 pages).
 - [2295] [Towards Energetic Quantum Advantage in Trapped-Ion Quantum Computation](#), by Francisca Góis, Marco Pezzutto, and Yasser Omar, arXiv, April 2024 (10 pages).
 - [2296] [Energetic advantages for quantum agents in online execution of complex strategies](#), by Jayne Thompson, Paul M. Riechers, Andrew J. P. Garner, Thomas J. Elliott, and Mile Gu, arXiv, March 2025 (5 pages).
 - [2297] [Towards identifying possible fault-tolerant advantage of quantum linear system algorithms in terms of space, time and energy](#), by Yue Tu, Junyu Liu, et al., arXiv, February 2025 (28 pages).
 - [2298] [Accelerating the drive towards energy-efficient generative AI with quantum computing algorithms](#), by Frederik F. Flöther, Jan Mikolon, and Maria Longobardi, arXiv, August 2025 (11 pages).
 - [2299] [Gradient Flow Decoding for LDPC Codes](#), by Tadashi Wadayama, Kensho Nakajima, and Ayano Nakai-Kasai, arXiv, March 2023 (6 pages).
 - [2300] [Competition of decoherence and quantum speed limits for quantum-gate fidelity in the Jaynes-Cummings model](#), by Sagar Silva Pratapsi, Lorenzo Buffoni, and Stefano Gherardini, arXiv, March 2024 (14 pages).
 - [2301] [Designing Energy-Efficient Quantum Computers Through Prediction and Reduction of Cooling Requirements for Cryogenic Electronics](#), by Michael Martin et al., NREL, November 2020.
 - [2302] [Carbon Emissions of Quantum Circuit Simulation: More than You Would Think](#), by Jinyang Li, Qiang Guan, Dingwen Tao, and Weiwen Jiang, arXiv, July 2023 (3 pages).
 - [2303] [Scaling up to Problem Sizes: An Environmental Life Cycle Assessment of Quantum Computing](#), by Sylvain Cordier, Karl Thibault, Marie-Luc Arpin, and Mourad Ben Amor, arXiv, March 2025 (26 pages).
 - [2304] [Energetics of a Single Qubit Gate](#), by Jeremy Stevens, Andrew Jordan, Audrey Bienfait, Alexia Auffèves, Benjamin Huard, et al., Physical Review Letters, September 2022  (preprint on [arXiv](https://arxiv.org/abs/2209.00000)).
 - [2305] [Calorimetry of a phase slip in a Josephson junction](#), by E. Gümüş, J. P. Pekola, H. Courtois, W. Belzig, C. B. Winkelmann, et al., Nature Physics, January 2023  (preprint on [HAL](https://arxiv.org/abs/2301.00000)).
 - [2306] [Quantum refrigeration powered by noise in a superconducting circuit](#), by Simon Sundelin, Mohammed Ali Aamir, Vyom Manish Kulkarni, Claudia Castillo-Moreno, and Simone Gasparinetti, Nature Communications, January 2026 (8 pages).
 - [2307] [Energetic cost of measurements using quantum, coherent, and thermal light](#), by Xiayu Linpeng et al., PRL, June 2022 (13 pages)  (preprint on [arXiv](https://arxiv.org/abs/2206.00000)).
 - [2308] [Quantum energetics of a non-commuting measurement](#), by Xiayu Linpeng, Nicolò Piccione, Maria Maffei, Léa Bresque, Samyak P. Prasad, Andrew N. Jordan, Alexia Auffèves, and Kater W. Murch, arXiv, November 2023 (9 pages).
 - [2309] [Fundamental mechanisms of energy exchanges in autonomous measurements based on dispersive qubit-light interaction](#), by Nicolò Piccione, Maria Maffei, Xiayu Linpeng, Andrew N. Jordan, Kater W. Murch, and Alexia Auffèves, arXiv, July 2024 (9 pages).
 - [2310] [Heat and Work in Quantum Thermodynamics: a Cybernetic Approach](#), by William Rupush and Oscar Grånäs, arXiv, March 2024 (9 pages).
 - [2311] [Quantum thermodynamic methods to purify a qubit on a quantum processing unit](#), by Andrea Solfanelli, Alessandro Santini, and Michele Campisi, arXiv, March 2022 (7 pages).
 - [2312] [Quantum computing is a path to energy-efficient AI](#), by Alan Baratz, FastCompany, June 2025.
 - [2313] [Thermodynamic significance of QUBO encoding on quantum annealers](#), by Emery Doucet, Zakaria Mzaouali, Reece Robertson, Bartłomiej Gardas, Sebastian Deffner, and Krzysztof Domino, arXiv, January 2026 (17 pages).
 - [2314] [Quantum annealers as programmatic thermal machines](#), by Jakub Pawłowski, Tomasz Śmierchalski, Fengping Jin, Bartłomiej Gardas, Sebastian Deffner, and Zakaria Mzaouali, arXiv, August 2026 (21 pages).
 - [2315] [Blockchain with proof of quantum work](#), by Mohammad H. Amin, Andrew D. King, Samuel Kortas, et al., Physical Review A, March 2026 (24 pages)  (preprint on [arXiv](https://arxiv.org/abs/2603.00000)).
 - [2316] [Quantum dynamics for energetic advantage in a charge-based classical full-adder](#), by João P. Moutinho, Marco Pezzutto, Sagar Pratapsi, Francisco Ferreira da Silva, Silvano De Franceschi, Sougato Bose, António T. Costa, and Yasser Omar, PRX Energy, July 2023 (11 pages).
 - [2317] [Energy-efficient entanglement generation and readout in a spin-photon interface](#), by Maria Maffei, Loïc Lanco, Andrew Jordan, Alexia Auffèves, et al., Quantum Journal, May 2022 (6 pages).
 - [2318] [Classical Half-Adder using Trapped-ion Quantum Bits: Towards Energy-efficient Computation](#), by Sagar Silva Prat-

- apsi, Patrick H. Huber, Patrick Barthel, Sougato Bose, Christof Wunderlich, and Yasser Omar, arXiv, April 2024 (10 pages).
- [2319] [Resource assessment of classical and quantum hardware for post-quench dynamics](#), by Joseph Vovrosh, Tiago Mendes-Santos, Hadriel Mamann, Kemal Bidzhiev, Fergus Hayes, Bruno Ximenez, Lucas Béguin, Constantin Dalyac, and Alexandre Dauphin, arXiv, November 2025 (15 pages).
- [2320] [A tweezer array with 6100 highly coherent atomic qubits](#), by Hannah J. Manetsch, Manuel Endres, et al., Nature, September 2025 (25 pages).
- [2321] [Trapping 11,000 Atoms in a Tweezer Array Generated by a Single Metasurface](#), by Yuqing Wang, Hui Zhai, et al., arXiv, June 2026 (6 pages).
- [2322] [PaQit: Energy-Runtime-Fidelity Co-Optimization for Neutral Atom Quantum Computers](#), by Jason Ludmir and Tirthak Patel, arXiv, August 2026 (11 pages).
- [2323] [Energetics of Rydberg-atom Quantum Computing](#), by Óscar Alves, Marco Pezzutto, and Yasser Omar, arXiv, January 2026 (12 pages) The QFT/DFT comparable in this paper is plagued by not accounting for the cost of encoding the data in the quantum register before the QFT is executed.
- [2324] [Minimizing resource overhead in fusion-based quantum computation using hybrid spin-photon devices](#), by Stephen C. Wein, Pascale Senellart, Boris Bourdoncle, Shane Mansfield, et al., PRX Quantum, December 2025 (22 pages).
- [2325] [Monitoring the generation of photonic linear cluster states with partial measurements](#), by Valentin Guichard, Pascale Senellart, et al., arXiv, May 2025 (10 pages).
- [2326] [Practical blueprint for low-depth photonic quantum computing with quantum dots](#), by Ming Lai Chan, Aliko Anna Capatos, Peter Lodahl, Anders Søndberg Sørensen, and Stefano Paesani, arXiv, July 2025 (23 pages).
- [2327] [Quantum Energetic Advantage before Computational Advantage in Boson Sampling](#), by Ariane Soret, Nessim Dridi, Stephen C. Wein, Valérian Giesz, Shane Mansfield, and Pierre-Emmanuel Emeriau, arXiv, January 2026 (13 pages).
- [2328] [Optimizing Wigner Negativity in Scattering Processes Using Energetic Cost Functions](#), by Kian Hwee Lim, Kiarn T. Laverick, Sahil Sardar Jafar, Samyak P. Prasad, Maria Maffei, and Alexia Auffèves, arXiv, June 2026 (5 pages).
- [2329] [Hamiltonian Benchmark of a Solid-State Spin-Photon Interface for Computation](#), by Tejas Acharya, Loïc Lanco, Olivier Krebs, Hui Khoon Ng, Alexia Auffèves, and Maria Maffei, arXiv, February 2026 (10 pages).
- [2330] [Mitigating photon loss in linear optical quantum circuits](#), by James Mills and Rawad Mezher, arXiv, March 2026 (30 pages).
- [2331] [Thermodynamics of autonomous optical Bloch equations](#), by Samyak Pratyush Prasad, Maria Maffei, Patrice A. Camati, Cyril Elouard, and Alexia Auffèves, arXiv, July 2025 (19 pages).
- [2332] [Accessing which-path information in the absorption and emission of light by a quantum dot in a Ramsey sequence](#), by I. Maillette de Buy Wenniger, M. Maffei, S. C. Wein, S. P. Prasad, H. Lam, D. Fioretto, A. Lemaitre, I. Sagnes, C. Antón-Solanas, P. Senellart, and A. Auffèves, arXiv, 2026 (12 pages).
- [2333] [An Energetic Constraint for Qubit-Qubit Entanglement](#), by Kiarn T. Laverick, Samyak P. Prasad, Pascale Senellart, Maria Maffei, and Alexia Auffèves, arXiv, 2026 (9 pages).
- [2334] [Estimating the electrical energy cost of performing arbitrary state preparation using qubits and qudits in integrated photonic circuits](#), by Maria Carolina Volpato, Gabriel da Silva Sampaio, and Pierre-Louis de Assis, arXiv, February 2026 (24 pages).
- [2335] [Thermodynamics of Quantum Reservoir Computing](#), by Lixiang Ding and Xingze Qiu, arXiv, July 2026 (26 pages).
- [2336] [Energetic Cost of Temporal Information Processing in Quantum Reservoirs](#), by Gabriele Cenedese, Gonzalo Manzano, Gian Luca Giorgi, and Roberta Zambrini, arXiv, August 2026 (17 pages).
- [2337] [Sustainable Quantum Computing: Opportunities and Challenges of Benchmarking Carbon in the Quantum Computing Lifecycle](#), by Nivedita Arora and Prem Kumar, arXiv, August 2024 (5 pages).
- [2338] [Frugal Machine Learning](#), by Mikhail Evchenko, Joaquin Vanschoren, Holger H. Hoos, Marc Schoenauer, and Michèle Sebag, arXiv, November 2021 (31 pages).
- [2339] [Energy-Efficient Training and Inference in Large Language Models: Optimizing Computational and Energy Costs](#), by Krishnam Raju Narsepalle, International Journal of Computer Applications, June 2025 (13 pages).
- [2340] [Efficient Large Language Models: A Survey](#), by Zhongwei Wan, Mi Zhang, et al., arXiv, May 2024 (67 pages).
- [2341] [EfficientLLM: Efficiency in Large Language Models](#), by Zhengqing Yuan, Yanfang Ye, et al., arXiv, May 2025 (85 pages).
- [2342] [A Survey on the Optimization of Large Language Model-based Agents](#), by Shangheng Du, Jiabao Zhao, Jinxin Shi, Zhentao Xie, Xin Jiang, Yanhong Bai, and Liang He, arXiv, March 2025 (42 pages).
- [2343] [K2-Think: A Parameter-Efficient Reasoning System](#), by Zhoujun Cheng, Eric P. Xing, et al., arXiv, September 2025 (23 pages).
- [2344] [SpikingBrain Technical Report: Spiking Brain-inspired Large Models](#), by Yuqi Pan, Guoqi Li, et al., arXiv, September 2025 (30 pages).
- [2345] [Jet-Nemotron: Efficient Language Model with Post Neural Architecture Search](#), by Yuxian Gu, Han Cai, et al., arXiv, August 2025 (20 pages).
- [2346] [Measuring the environmental impact of AI inference](#), by Amin Vahdat and Jeff Dean, Google Cloud Blog, August 2025.
- [2347] [Modeling PFAS in Semiconductor Manufacturing to Quantify Trade-offs in Energy Efficiency and Environmental Impact of Computing Systems](#), by Mariam Elgamal, Abdulrahman Mahmoud, Gu-Yeon Wei, David Brooks, and Gage Hills, arXiv, May 2025 (7 pages).
- [2348] [Engineering cryogenic setups for 100-qubit scale superconducting circuit systems](#), by S. Krinner et al., EPJ-QT, May 2019 (29 pages).
- [2349] [CMOS-based cryogenic control of silicon quantum circuits](#), by Xiao Xue, Bishnu Patra, Jeroen P. G. van Dijk, Nodar Samkharadze, Andrea Corna, Brian Paquelet Wuetz, Amir Sammak, Giordano Scappucci, Menno Veldhorst, Fabio Sebastianio, Masoud Babaie, Edoardo Charbon, and Lieven M. K. Vandersypen, Nature, May 2021 (6 pages).
- [2350] [EU Strategic Research and Industry Agenda 2030](#), by European Union Flagship, EU Flagship, February 2024 (270 pages).
- [2351] [Quantum Energy Initiative website](#), by Olivier Ezratty et al, QEI, August 2022.
- [2352] [Quantum technologies need a Quantum Energy Initiative](#), by Alexia Auffèves, PRX Quantum, June 2022.
- [2353] [Are quantum computers really energy efficient?](#), by Sophia Chen, Nature Computational Science, June 2023.
- [2354] [P3329: Standard for Quantum Computing Energy Efficiency](#), IEEE, 2023.
- [2355] [Green HPC: An analysis of the domain based on Top500](#), by Abdessalam Benhari, Denis Trystram, Fanny Dufossé, Yves Denneulin, and Frédéric Desprez, arXiv, March 2024 (15 pages).
- [2356] [Operating Fiber Networks in the Quantum Limit](#), by Janis Nötzel and Matteo Rosati, arXiv, May 2022 (30 pages).
- [2357] [Energetic Analysis of Emerging Quantum Communication Protocols](#), by Raja Yehia, Yoann Piétri, Carlos Pascual-

García, Pascal Lefebvre, and Federico Centrone, arXiv, March 2025 (29 pages).

[2358] [Quantification of the energy consumption of entanglement distribution](#), by Karol Horodecki, Marek Winczewski, Leonard Sikorski, Paweł Mazurek, Mikołaj Czechlewski, and Raja Yehia, arXiv, July 2025 (46 pages).

[2359] [Quantum computing in the net-zero transition: energy production, management, and efficiency](#), by Irais Bautista,

Stefano Lodi, Yasser Omar, Olivier Ezratty, et al., EERA website, May 2026 (30 pages).

[2360] [Quantum Technologies in the Context of Climate Change: Emphasizing Sustainability in a Responsible Innovation Approach to Quantum Innovation - Ethics and Society](#), by Dakota Root, NanoEthics, March 2025 (15 pages).

This **Energetics** subsection contains 234 bibliographical references and notes containing at least 3,910 pages.

5.5 Lieb-Robinson bounds

The story of this new subsection is the following. I happened to meet, in January 2025, Lucas Leclerc, a quantum computing physicist working at Pasqal, who told me he was interested in the effects of the so-called Lieb-Robinson bound. It was like discovering a new Pandora’s box. I was struck by its potential implications for scaling quantum computers. I repeatedly asked questions about it to a number of quantum physicists in both academic and industry vendor circles. Some knew about it, most had fuzzy ideas about its implications, or they had quite different views on its practical effects. I wanted to add something about it in this book and here we are. This part is written on a relatively unexplored domain when dealing with gate-based quantum computing, so all claims and ideas are food for thought more than established science.

The Lieb-Robinson (LR) bound was proved in 1972 by Elliott H. Lieb and Derek W. Robinson who were then affiliated with the MIT and CNRS / Université Aix Marseille in France [2361]. It sets a fundamental limit on the speed at which information and correlations can propagate in quantum many-body systems with local interactions. It defines an effective “light cone” beyond which correlations decay exponentially (Figure 303). An analogous but distinct light-cone structure exists for entanglement itself. The LR bound constrains commutators of observables, which does not directly control an entanglement measure, so a separate result is needed. Sigal established one in 2026 for bipartite systems with localized couplings, yielding a lower bound on the time needed to transport entanglement to a distant node of a quantum network or to maintain it there [2362].

This constraint may impact quantum computing architectures that rely on spatially distributed qubits and either local quantum gates in gate-based modes or even analog quantum simulations or annealing. As quantum processors scale up, LR bounds could become a bottleneck for entanglement distribution, circuit depth, and algorithmic efficiency. There are several variations of this bound depending on the type of system, particularly considering long range interactions [2363–2365].

The literature uses the terms “LR bound” or “LR bounds” depending on whether it refers to a singular result or a broader class of inequalities. An “LR bound” is about a specific formulation of the Lieb-Robinson bound in a given quantum system. It is used when discussing a particular system, like a qubit type. In contrast, “LR bounds” is about LR bounds in general applied to different quantum systems with their own specificities like interactions and dimensions. This plural form is appropriate when comparing multiple versions or highlighting the evolution of the theory. Using them depends on whether you are discussing a single instance or looking at the concept in general.

LR bounds were first experimentally observed in a quantum many-body system in 2012 [2366]. By quenching a

one-dimensional Bose gas in an optical lattice, meaning generating a rapid change in the system’s conditions, the experimenters tracked how correlations spread over time and revealed a light-cone-like structure in the dynamics. The propagation was attributed to quasiparticle pairs moving with finite velocity, confirming the existence of an effective causal horizon. The experiment demonstrated that correlation spread in a quantum simulation is bounded, enabling predictable entanglement growth. Quench dynamics can be used to probe system response and validate simulation fidelity. Theoretical papers analyze the effects of the LR limit and its related light-cone, like in non-Fermi liquids in condensed matter [2367].

Two 2014 trapped-ion experiments extended the picture to long-range interactions, one observing quasiparticle-mediated entanglement propagation with a well-defined light cone [2368] and the other observing correlation propagation that becomes effectively non-local when the interaction decay exponent is small enough [2369]. On gate-based hardware, Braumüller et al. measured out-of-time-ordered correlators on a 3×3 superconducting lattice and inferred the light cone from the OTOC data [2370], which remains the cleanest demonstration that the LR picture is measurable on a superconducting QPU.

As a consequence, the light-cone velocity becomes a design parameter for timing and control logic in both analog quantum simulators and gate-based quantum computers. While the LR bounds literature is focused on analog systems, it could have serious implications on the design of gate-based quantum computers.

Interestingly, this topic is at the crossroads of quantum physics, quantum computing, compiler design, quantum communications, and even quantum foundations and quantum thermodynamics (for this respect, see the contribution from Moallison F. Cavalcante et al. in the Roadmap on Quantum Thermodynamics published in 2025 [2226] and [2371]). This is what we uncover in this subsection, looking at the various potential implications of LR bounds on gate-based quantum computers and analog quantum simulators. This is a topic that hasn’t been much discussed in academic and commercial circles. One of the reasons being that the effects of LR bounds may have some impact only on large QPU settings that do not exist yet. But they are currently being designed, so it deserves to be considered.

5.5.1 Probing information propagation

In the context of Lieb-Robinson bounds, the term information refers to the ability of one part of a quantum system to influence or affect another part. More specifically, it captures how local observables such as spin orientation at one site can impact measurements at distant sites over time. This notion of information is tied to the spread of correlations, entanglement, or disturbances through a quantum lattice which can be an atom lattice like in an analog quantum simulator or a set of




Figure 303: An illustration of the Lieb-Robinson light-cone on a lattice of physical qubits in a gate-based model, showing the extended influence over time (t_0 , t_1 , t_2 and t_3) of a local modification on the center qubit in green on a grid of (gate-based) qubits that are already pre-entangled. Two situations can arise. One is the establishment step by step of a large entangled state. It will be based on a series of two-qubit gates connecting the qubits together in the grid it can start from a single point or from multiple points. At each step, the gate time will probably be greater than the Lieb-Robinson bound of the system. Entanglement does not make a local gate act physically on remote qubits, and applying a gate does not require waiting for its influence to reach the whole register. In a digital circuit, what a gate can affect is its causal cone, generated by the subsequent gates of the circuit itself, and legal parallelism follows from shared-qubit conflicts, commutation and hardware control constraints, all of which the compiler is supposed to track through dependency graphs. This figure can be read as a circuit causal cone. This is actually a question that is not well explored for gate-based quantum computing since existing circuits are still quite small. When they grow in size and quantum processing units become interconnected over some distance in the meters range, the LR bound may be at play. Information propagation on all the register qubits may take more time than the gate times that created the modification. It could potentially create new constraints for gate synchronization. A Lieb-Robinson wavefront in the strict sense belongs to continuous evolution under a specified local Hamiltonian, as in an analog simulator or an always-on interaction model, and not to a sequence of discrete gates. Where real waiting times appear in a large machine, they come from link latency, heralded entanglement generation, feed-forward and clock distribution, which are engineering causal constraints to be modeled on their own terms. The light blue corresponds to the propagation wall with vanishing propagation effects. Grey qubits are the ones not yet affected by the modification of the green qubit. © Olivier Ezratty, September 2025-September 2026.

solid state qubits on a chip like with superconducting and silicon spin qubits.

When a quantum system is modified locally, for example by flipping a spin at a specific site, or with applying a quantum gate to one or two qubits, the information about that change propagates outward over time through the system like a ripple effect. At least, when the system is already in an entangled state. With gate-based systems, a global entanglement is established step by step with quantum gates progressively reaching out more and more qubits in the register.

The Lieb-Robinson bound provides a mathematical framework that describes this propagation phenomenon, showing that it is not instantaneous. Instead, it is limited by a finite speed, known as the Lieb-Robinson velocity (v_{LR}).

This propagation of information is expressed through the commutator between two operators, typically denoted as $[A_x(t), B_y]$, where $A_x(t)$ is an observable evolved in time at site x (the qubit you “touch”), and B_y is another observable at site y (a remote qubit that is for example already entangled with the source qubit due to previous quantum gate operations and that you measure). The magnitude of this commutator quantifies how much the two observables influence each other. A commutator equal to zero means that the two systems evolve inde-

pendently and don’t influence each other. When it is non zero, it means one influences the other.

The Lieb-Robinson bound says:

$$\|[A_x(t), B_y]\| \leq C \|A\| \|B\| \min(|X|, |Y|) \cdot e^{-\frac{d(x,y)-vt}{\xi}}$$

where $d(x,y)$ is the distance between quantum sites, v is the Lieb-Robinson velocity, ξ is a decay length, t is time, $|X|$ and $|Y|$ are the numbers of sites supporting A and B , and C is a prefactor constant depending on multiple factors (lattice geometry and connectivity, local Hilbert space dimension, interaction strengths and range, decay properties of the interaction). The bound is therefore relative to the size of the observables, not an absolute number of bits or an amplitude.

The double bars around the A, B commutator correspond to operator norms. The norm of an operator O is $\|O\| = \sup_{\|\psi\|=1} \|O\psi\|$, meaning, when taking all possible unit vectors ψ , apply the operator O , and find the largest resulting vector length. It is a measure of how strong or large the operator’s action can be on a given vector. This norm gives a quantitative bound on how A influences B .

The inequality tells us that outside the light cone defined by $vt < d(x,y)$, the commutator is exponentially small.

What this bounds is the sensitivity of the observable B_y to a perturbation generated by A_x , at that time and for those observables. A vanishing commutator does not by itself mean that the two subsystems evolve independently or that they cannot influence each other, it means that these particular observables commute at that time.

The Lieb-Robinson bound demonstrates that this commutator is exponentially small outside a causal region labeled a “light cone” and defined by the Lieb-Robinson velocity. Thus, information spreads through the system in a wave-like manner, respecting a form of causality, even in non-relativistic quantum systems.

Recent papers related to the LR bounds are even tightening it, meaning that they consider that older LR bounds gave a very loose upper limit on how fast information could spread in a quantum system [2372, 2373]. In [2374] published in 2020, Zhiyuan Wang and Kaden R.A. Hazzard tightened the LR bound, meaning that the maximum speed at which information or entanglement can propagate is lower than previously thought. The commutator between distant operators are decaying faster, particularly in systems with large local Hilbert spaces like with qudits (instead of qubits). Their improvement is quantitative: for a Heisenberg model with spin S they obtain a velocity bounded by a constant instead of the previous $v \propto S$, and for multiorbital Hubbard models with N orbitals they get $v \propto \sqrt{N}$ instead of $v \propto N$. Even for spin-1/2 Ising and Fermi-Hubbard models the velocity improves by an order of magnitude. For long-range interactions, a parallel line of work established nearly linear [2375], then strictly linear light cones [2376, 2377] and finite scrambling speed [2378].

This order-of-magnitude sensitivity is worth keeping in mind for the rest of this subsection. LR bounds are upper bounds on a velocity, and the constant prefactors are notoriously loose. A numerical value for v_{LR} quoted for a given hardware platform should be read as “no faster than”, not as “this is how fast information actually travels”. Any engineering conclusion that depends on the numerical value rather than on its scaling is fragile.

This concept provides a quantum equivalent of causality, ensuring that distant parts of a system cannot influence each other faster than a certain speed. It supports the theoretical foundations of quantum simulation, quantum error correction, and entanglement growth. It also plays a role in proving area laws for entanglement entropy, the stability of quantum phases, and the design of efficient algorithms in quantum computing.

5.5.2 Nonlocality vs information propagation

Nonlocality refers to measurement correlations that violate classical physics, such as those observed in Bell inequality violation experiments. These correlations are incompatible with the existence of local hidden variables, but they do not allow faster-than-light signaling. Nonlo-

cal correlations are instantaneous but they don’t enable the transmission of some information. The randomness of outcomes ensures this. Spacelike separated parties cannot directly use correlations to transmit controllable information instantaneously. The teleportation protocol which transmits a quantum state from one quantum object to another is limited by the speed of its mandatory classical channels.

Lieb-Robinson bounds are applicable in a different context, the evolution of local observables in systems with short-range interactions. They don’t rule out entanglement but constrain how fast observable effects of entanglement can spread. Nonlocality is a feature of quantum correlations under measurement while Lieb-Robinson bounds apply to unitary time evolution of operators and the growth of their support [2379].

Several mathematical tools describe the evolution of many-body quantum systems subject to LR bounds, including quasi-local algebras [2380], interaction norms [2365, 2381], Heisenberg evolution [2382], commutator estimates [2383], quasi-adiabatic continuation [2384], and generalizations to long-range interactions [2363]. For the sake of accessibility, we won’t dig much on this here. It is also beyond my knowledge grade!

Lieb-Robinson bounds relate to dynamics generated by local Hamiltonians on a lattice. This is fully compatible with Bell nonlocality. LR bounds constrain how fast influences produced by local interactions can propagate under unitary time evolution while Bell nonlocality concerns static measurement correlations in an already-prepared quantum state and is not enabling superluminal signaling.

A Lieb-Robinson light-cone effect was one of the by-products of the famous Google OTOC experiment published in October 2025 in Nature [2385]. It was observed when running an algorithm using interferences on 65 qubits of a 103-qubit operational variant of the 105-qubit Willow processor (Figure 305).

5.5.3 Vector state vs quantum channels

In textbook quantum computing, a quantum register evolves via unitary gates. Each gate acts on a subset of qubits, and the full evolution is described by tensoring the gate unitary with identities for untouched qubits. The mathematical gate is a unitary matrix multiplied with the register vector state, which creates a new vector state as shown below with actual values on a 4-qubit register.

Let’s look at the idealized mathematics model behind it with considering a quantum register of N qubits. Its corresponding Hilbert space is:

$$\mathcal{H} = \bigotimes_{i=1}^N \mathbb{C}^2$$

The qubit register is usually described by a pure state which is a vector $|\psi\rangle \in \mathcal{H}$ of dimension 2^N . Its evolution depends on the gates we apply consecutively to the register. Suppose we apply a quantum gate U_{ij} to qubits i and j . The full unitary that is applied on the register is a matrix of dimension $2^N \times 2^N$:

$$\tilde{U}_{ij} = I_1 \otimes \cdots \otimes I_{i-1} \otimes U_{ij} \otimes I_{j+1} \otimes \cdots \otimes I_N, \quad j = i+1$$

For non-adjacent qubits, the gate acts on positions i and j only, and the identities on all sites between them are kept, which is written more compactly as $\tilde{U}_{ij} = U_{ij} \otimes I_{\{1..N\} \setminus \{i,j\}}$ up to a reordering of the tensor factors.

After the quantum gate is applied, the register quantum state becomes the result of applying the unitary matrix $\tilde{U}_{ij}$ on the vector state $|\psi(0)\rangle$:

$$|\psi(t)\rangle = \tilde{U}_{ij}|\psi(0)\rangle$$

This assumes everything happens instantly, with no delay. It's the idealized version we learn when programming gate-based quantum computers. For example, when a CNOT gate is applied to a 4-qubit register like in the example below, the whole system is expected to update immediately. But this model is unrealistic and often confuses developers—especially when working with quantum emulators.

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \\ \alpha_5 \\ \alpha_6 \\ \alpha_7 \\ \alpha_8 \\ \alpha_9 \\ \alpha_{10} \\ \alpha_{11} \\ \alpha_{12} \\ \alpha_{13} \\ \alpha_{14} \\ \alpha_{15} \end{bmatrix} = \begin{bmatrix} \alpha_0 \\ \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_5 \\ \alpha_4 \\ \alpha_7 \\ \alpha_6 \\ \alpha_8 \\ \alpha_9 \\ \alpha_{10} \\ \alpha_{11} \\ \alpha_{13} \\ \alpha_{12} \\ \alpha_{15} \\ \alpha_{14} \end{bmatrix}$$

In real life, things are way different (Table 13). Gates are not applied instantaneously but are quantum evolutions over time governed by a time-dependent Hamiltonian $H(t)$, and their effect on the quantum register is more complicated than just applying a simple unitary to a vector state.

Instead of discrete gates, you must consider a local Hamiltonian:

$$H = \sum_{\langle i,j \rangle} h_{ij}$$

where each h_{ij} acts on neighboring qubits. The time evolution of the quantum system is governed by:

$$|\psi(t)\rangle = U(t)|\psi(0)\rangle, \quad \text{with} \quad U(t) = e^{-iHt}$$

Due to the LR bound, the influence of h_{ij} on distant qubits is not instantaneous. The bound is:

$$\|[A_x(t), B_y]\| \leq C \cdot e^{-\frac{d(x,y)-vt}{\xi}}$$

with $d(x,y)$ being the distance between the qubits, v the Lieb-Robinson velocity, and ξ is a decay length.

To describe local evolution with delay, you need to use a quantum channel $\mathcal{E}_t$, which is a completely positive trace-preserving (CPTP) map.

We covered it in the “Gate based” section of this book. It is a $4^N \times 4^N$ matrix, so 16^N entries, N being the number of qubits:

$$\rho(t) = \mathcal{E}_t(\rho(0))$$

This channel can be written using Kraus operators $\{K_\alpha(t)\}$ (which satisfy $\sum_i K_i^\dagger K_i = I$):

$$\mathcal{E}_t(\rho) = \sum_i K_i(t) \rho K_i^\dagger(t)$$

The individual Kraus operators are not required to be local, and the Kraus decomposition is not unique. What the LR bound constrains is the channel itself, which is quasi-local: $\mathcal{E}_t$ can be approximated, with an error decaying exponentially outside the light cone, by a channel supported only on a neighborhood of radius $v_{LR}t$. David Poulin proved the corresponding bound for general Markovian dynamics in 2010 [2386].

In summary, instead of having an instant unitary applied to our qubit register vector state, there is a continuous time evolution of our quantum system described by its density matrix and a quantum channel acting on it over some time. The end-result is the same, but we now have here a better understanding on how this happens over time.

Two caveats have to be mentioned at this point, because they condition everything that follows.

First, a local unitary applied to qubit A can never change any local observable measured on qubit B, however entangled A and B are and however long you wait. This is due to the no-signaling property, and it is exact, not asymptotic: $\text{Tr}_A[(U_A \otimes I)\rho(U_A^\dagger \otimes I)] = \text{Tr}_A[\rho]$. So the question “how long before flipping A becomes visible

at B” has no answer in terms of a delay. The change appears only in joint correlators, and only if a coupling Hamiltonian connects A and B. LR bounds describe the growth of the support of a Heisenberg-evolved operator under that coupling, not the arrival of a signal.

Second, in a digital gate-based machine that coupling is the gate sequence itself. Between gates, tunable couplers are ideally off and the residual always-on Hamiltonian is a parasitic term, not the computational one. The light cone of a depth-D circuit with nearest-neighbor two-qubit gates therefore expands by one lattice site per layer, giving an effective velocity $v_{\text{circuit}} \sim a/t_{\text{gate}}$. With $a = 200 \mu\text{m}$ and $t_{\text{gate}} = 100 \text{ ns}$ this is about $2 \times 10^3 \text{ m/s}$, far below any free-evolution v_{LR} quoted for the same hardware. The binding constraint is the circuit light cone, and compilers already track it exactly through the circuit dependency graph. There is no additional physical delay to wait out, and no “wait” instruction to insert.

Where the LR bound impacts gate-based computing is elsewhere: on always-on residual couplings and crosstalk between idle qubits, on analog blocks embedded in digital circuits such as Hamiltonian simulation or Rydberg global drives, on the physical latency of inter-processor links, which is a propagation delay rather than an LR effect, and on the classical control and feedforward latency. Those are the cases developed in the rest of this subsection.

Now, let’s visualize this propagation with a toy lattice of qubits that represent at a small scale how it would work on static qubits with nearest neighbor connectivity. This propagation is at play in several situations as described in Figure 304:

- Entanglement establishment between many quantum objects when you start an analog simulation with Rydberg state cold atoms [2382].
- When modifying a qubit in an already large entangled set of qubits in the case of a gate-based quantum computer [2387]. The propagation of the modification to the large entangled state will be progressive (phase 3 in Figure 304).
- This will also impact the establishment of entanglement as the entangled state grows. Let’s say for example that we apply a series of CNOT to entangle N qubits. At the beginning, the LR bound will have no effect, the size of the entangled state being small and the first entangled qubits within the LR light cone. But as it grows, its extension with additional CNOT will hit the LR wall when the distance between the first entangled qubit and the new entangled qubits is outside the LR light cone. We may need to wait till the information propagation from the new CNOT impacts the first qubit (phase 2 in Figure 304).

But in phase 1 in Figure 304, local operations do not propagate information to other qubits and the LR bound is not at play.

5.5.4 Gate-based quantum computing

We will look here at the impact of LR bounds on the implementation of practical gate-based quantum computing at scale. It will depend on the qubit type and their spread in space, their interconnect and will have an influence on how to design quantum circuits and compilers.

The speed of influence propagation depends on several hardware and model-specific factors:

- **Geometry and degree of connectivity z** with a higher connectivity might accelerate information propagation.
- **Interaction strength J** : LR velocity scales roughly with J . The LR velocity scales as $v_{\text{LR}} \sim e z J a / \hbar$, where a is the lattice spacing and e the Euler number coming out of the standard proof.
- **Interaction range and decay α** : finite-range interactions yield exponential cones, while power-law $1/r^\alpha$ interactions give polynomial tails.
- **Qubit type and gate mechanism**: Different hardware platforms yield drastically different LR behavior.

Qubit type specifics

With static qubits like superconducting qubits, spin qubits, cavity qubits, cold atoms and trapped ion qubits, gate-based quantum computers typically rely on nearest-neighbor interactions. The LR velocity v_{LR} depends on the interaction strength and geometry between the qubits (see a summary in Table 14).

- **Superconducting qubits.** Superconducting qubits have a Lieb-Robinson velocity in the 10^4 to 10^5 m/s range. The lower end follows from a residual coupling $J/\hbar \sim 10 \text{ MHz}$ at a pitch $a = 200 \mu\text{m}$, giving $v_{\text{LR}} \sim 2\pi J a / \hbar \approx 1.3 \times 10^4 \text{ m/s}$. Propagating a perturbation across a 1D array of 100 qubits spaced $200 \mu\text{m}$ apart, so a 2 cm chip, then takes 200 ns at 10^5 m/s and $2 \mu\text{s}$ at 10^4 m/s . Both are shorter than transmon coherence times, and the first is comparable to gate times of 20 ns to 400 ns. For large-scale QPUs with spatially distributed qubits, gate operations act as local updates to the quantum register state, with negligible immediate impact on remote subsystems.

The next estimate, for multiple processors connected by microwave flexible cables, needs three corrections before it can be used, and I keep it here mostly to show what the LR formalism does and does not say about a modular architecture.




Figure 304: Understanding the effect of the LR bound on a gate-based quantum circuit with many qubits. Phase 1 contains only single qubit gates which are acting locally on each qubit. They don't propagate information to other qubits so the LR bound is not at play. Phase 2 contains entangling CNOT gates, here, with additional single qubit rotations, which progressively establishes entanglement with an increasing number of qubits. If we use a two-qubit gate time t_1 as a reference and for a large number of qubits, at t_1 after the gates on qubit 10 are done, the information propagation will have not yet reached the first entangled qubit. In phase 3, local operations impact all qubits which are already entangled. It can change correlations between the qubits but not their observables due to quantum no-signaling. Likewise, at time t_2 which is usually shorter than time t_1 , information propagation will not have reached all qubits. And this is for just one of the gates of the four gate cycles in the circuit. Can these gates be parallelized or do we need to wait till all information is propagated before executing the next gate? The qubit numbers used here (1, 10, 23) are small, just for the sake of explanation. In real-life systems, you will need to use a much larger number of physical qubits spread on multiple processing interconnected units to see the potential detrimental effect of LR bounds. © Olivier Ezratty, September 2025-September 2026 for the caption.




Figure 305: In the Google OTOC experiment, this figure shows the information front and the light cone. The data shows an “information front” or wavefront which moves outward as the circuit time (t) increases. Information spreads via a biased diffusion process generated by the iSWAP gate used in the circuits. This process has an average velocity slower than the light cone. This leads to wavefront broadening, distinguishing it from “rapid-scrambling patterns” where the operator spreads ballistically at the velocity of the lightcone. While the exact speed of the light cone propagation (the maximum possible speed) is not explicitly stated in units related to the gate length (maximum possible qubits traversed per cycle), the discussion confirms that the average operational spread of entanglement in the chosen “slow-scrambling pattern” is deliberately slower than this maximum theoretical light-cone velocity. Source: [2385].

- Gate duration typical for single qubit gates: $\tau = 20$ ns.
- Module separation of $D = 2$ m, an assumption on the physical footprint of a multi-module system in the spirit of what IBM is planning with its Starling

Feature	Idealized gate model	Realistic LR-bound model
Evolution	Instantaneous unitary.	Time-evolved via local Hamiltonian.
Scope of a local gate	A local unitary tensored with identities acts locally, even on an entangled register: it changes the global state coordinates without sending any controllable signal to remote qubits.	Under continuous evolution generated by a local Hamiltonian, the support of an evolved observable grows at a bounded rate, which is what the light cone describes.
Mathematical form	$\tilde{U} \psi\rangle$	$U(t) = e^{-iHt}$, or $\mathcal{E}_t(\rho)$
Locality	Explicit via tensoring with identity.	Emergent via LR-bound and channel locality.
Decoherence and noise	Not taken into account.	Modeled with open-system CPTP evolution, which describes general noisy dynamics and is a third model, distinct from the first two and not a description of spatial delay as such.

Table 13: Three models that are often aggregated and are kept apart here: the abstract circuit model with its own notion of locality, closed-system evolution under a local Hamiltonian, to which Lieb-Robinson bounds apply, and open-system CPTP evolution, which models noise. A quantum channel is not required in order to describe a local operation with a delay, and a local unitary is already local in the first model. Compiled by Olivier Ezratty from the references cited in this subsection.

QPU hosting 200 logical qubits by 2029. IBM has not published a module pitch, so this is my own estimate, not a vendor figure.

- The relevant velocity for a coaxial or waveguide link is the microwave group velocity, $v \approx 0.6c \approx 2 \times 10^8$ m/s, not an on-chip v_{LR} . A transmission line is not a lattice of locally coupled sites, so the LR bound is not the right tool for it. Two meters of cable is a 10 ns propagation delay, which is a classical latency budget, comparable to a single-qubit gate.
- The decay length ξ is set by the microscopic lattice, so $\xi \sim a \sim 200$ μm and $\mu = 1/\xi \sim 5 \times 10^3 \text{ m}^{-1}$, not 5 m^{-1} . With the larger value the exponent becomes of order -10^4 and the commutator bound is numerically zero. The original choice of $\mu = 5 \text{ m}^{-1}$, which amounts to a 20 cm decay length on a chip, is what made the estimate land on a readable 5×10^{-5} .

The conclusion to draw is the opposite of a bottleneck. Free-evolution LR propagation between two modules two meters apart is utterly negligible, and what actually costs time is the deliberate transfer: the cable delay, the entanglement distribution rate, and the classical control latency. And, per the no-signaling point made earlier, flipping a qubit at A is not detectable at B by a local measurement at any time, whether you wait or not.

- **Spin qubits in silicon** are somewhat similar to superconducting qubits, with slightly slower gate speeds and tighter LR bounds. Exchange coupling is short-ranged and strongly tunable, which puts silicon spin qubits well inside the standard nearest-neighbor LR framework. The related question of when a spin flip can be said to have arrived at a distant site, which is an arrival-time rather than a commutator question, is treated in [2388]. In practice the schedule of a silicon QPU is dominated by shuttling and routing latencies rather than by any LR limit.

- **Trapped ions** implement long-range phonon-mediated interactions allowing faster entanglement. Coherence is not the practical ceiling here, since Quantinuum Helios operates 98 $^{137}\text{Ba}^+$ ions with T_2 in the second range [2389]. Shuttling and gate times dominate instead. Gong et al. [2390] showed that locality persists for interactions decaying as $1/r^\alpha$ when $\alpha > D$, the spatial dimension, and ion chains with a tunable α between 0 and 3 can be placed on either side of that threshold, which is precisely what the 2014 experiments exploited [2368, 2369].
- **Cold atoms using cavity QED and Rydberg states** offer dynamically reconfigurable coupling graphs with enhanced connectivity which may relax some constraints imposed by the related LR bound.
- **Photonic qubits** have a different story since they are flying qubits with different operation modes. Models such as Fusion-Based Quantum Computing (FBQC) sidestep many of these restrictions. Here, entanglement isn't built by propagating interactions through static sites, it is prepared into cluster states. Computation proceeds via single-qubit measurements. This allows for high concurrency and non-local computation without physical adjacency. Measurement-based paradigms enable logical operations to proceed in parallel, limited only by detection latency and feedforward classical control, not by the LR velocity. But photonic qubits have their own bottlenecks with trading spatial LR-style constraints for stochastic resource preparation and heralding bottlenecks:
 - Cluster state generation is usually probabilistic and resource-intensive. Fusion gates usually require heralding and entanglement purification.
 - Losses in photonic integrated circuits (PICs), fibers, and routers degrade computing operations yield.
 - Photon detection, particularly for multiple photon detection (with PNRDs, photon number resolution

detectors) remains imperfect with dark counts, low efficiencies, and timing jitter affect scalability.

- Probabilistic fusion operations depending on classical orchestration of many fusion attempts, raising a lot of questions about the statistical resource overhead.
- Feedforward latency in active switching and classical control can become a computational rate-limiter.

These are however not fundamental bounds like the LR bounds and they don't seem to affect the space-time optimization of photonic quantum computing operations.

For algorithms requiring global entanglement like Shor's factoring algorithm, QFT, QPE and VQE, the practical ceiling on the number of qubits that can be usefully entangled within coherence time is set by gate times and two-qubit error rates, not by the LR bound. For superconducting systems with $T_2 \sim 100 \mu\text{s}$, a 200 ns two-qubit gate and a 10^{-3} error rate, the budget is a few thousand gates, which caps a 1D entangling chain at the order of 100 qubits and a 2D patch at roughly 30×30 . The LR light cone over the same 100 μs would reach several thousand sites, so it is one to two orders of magnitude away from being the binding constraint. LR bounds become the limiting factor only when the physical extent of the register grows faster than its gate budget, which is the modular and distributed case discussed below.

As a result, it seems that the LR bound is an indication that the mantra "survival of the fittest" could be appropriate for large scale quantum computers. The smaller and tighter form factors may be the ones that circumvent the best these LR bounds effects on computing time.

Compilers

As a result of the scheduling constraints, some compiler strategies may need to be considered, particularly for physically large quantum registers. It somewhat depends on the potential non-commutativity of parallelized gate operations and on the speed at which sequenced gates propagate information throughout the qubits that are already entangled.

It may impact gate scheduling and sequencing. Will compilers have to create some "wait" sequence between gates to account for the correlations propagation on all physical qubits? For a purely unitary circuit the answer is no, since the circuit light cone is already tracked by the dependency graph. When measurements and feedforward are allowed, the unitary light cone can be bypassed. IBM researchers demonstrated in 2024 a constant-depth long-range CNOT across up to 101 IBM qubits by feeding forward 99 mid-circuit measurement outcomes, with a fidelity that beats the unitary SWAP-chain implementation beyond a certain separation [2391].

They extended the experiment in 2026 to multiple state preparations [2392]. Friedman et al. from the University of Colorado provided in 2022 the matching bound, discussed further below, on how much measurement can help [2393].

In physically realistic quantum architectures, particularly those spread across meters of interconnected chips, information and entanglement will not propagate instantaneously. A quantum circuit compiler may need to account for this causal structure and insert delays or reorder instructions, particularly in cases where a gate depends on correlations formed by previous operations over distant regions, entangling operations span across hardware modules with finite communication bandwidth and when synchronization across spatial boundaries is required for phase estimation steps that involve some collective interference. This will therefore be algorithm structure dependent.

The constraints to consider are multiple:

- **Interaction-aware modeling** with tracking which gates act on which regions of the register and estimate dependency distances.
- **Velocity budgeting** with estimating the effective propagation velocity to infer the minimum latency $\tau \geq d/v_{\text{eff}}$ between causally linked gate sequences, where for a digital circuit $v_{\text{eff}} = a/t_{\text{gate}}$ and for an inter-module link it is the physical signal velocity.
- **Latency budgeting across modules** rather than generic guard bands. A purely unitary circuit needs no extra wait beyond its dependency graph, so what has to be modeled explicitly is link propagation, entanglement generation rate, buffering, clock distribution and feed-forward deadlines. An explicit wait belongs in the schedule when one of those engineering constraints requires it, not because of a Lieb-Robinson bound, which applies to a specified interacting Hamiltonian.
- **Dependency graph** maintenance with annotations for distance-based delays [2394].
- **Model causality** using gate duration rather than spatial delay for long-range or fast global gates (ions, Rydberg cold atoms).
- **Parallelism exploitation** with identifying spacelike-separated subgraphs and execute gates concurrently where LR bounds guarantee isolation between these subsystems. This can validate parallelism with operations outside the LR cone which are provably causally disconnected. But this could come at the expense of a broad quantum advantage.
- **Synchronization protocol** for inter-module entanglement like with photonic links or microwave QPU interconnect, with synchronizing based on link latency and interaction duration rather than idealized circuit structure.

Platform	Coupling Mechanism	Range/Decay	Gate Times	Compiler Implications
Superconducting	Capacitive or resonator exchange	Finite-range	10–200 ns	Inside module: parallelize safely; cross-chip: schedule entanglement via explicit link latency.
Silicon spin	Exchange, shuttle, resonators	Nearest-neighbor or routed	50 ns–1 μ s	Transport dominates schedule. LR constraints secondary to routing.
Trapped ions	Phonon-mediated spin–spin	Long-range $1/r^\alpha$, tunable α	10–300 μ s	Global entanglement with compiler avoiding mode crowding more than LR waits.
Neutral atoms	Rydberg blockade	Finite radius r_b , fast decay	50 ns–1 μ s	Tile operations by blockade radius. LR enables safe parallel groups.
Photons	Measurement based applied on cluster states and probabilistic fusion gates between clusters.	Becomes a problem with photon losses.	Depends on classical operations.	No LR light-cone. Scheduling depends on photon detection, feedforward, and classical control speed.

Table 14: Comparison of the impact of the Lieb-Robinson limit on various types of qubits. Each row corresponds to an interaction model, a hardware generation and a date that should be stated with its own primary reference, and the gate time ranges quoted are architecture-specific rather than modality-wide. For photons, the correct statement is not that there is no light cone, since relativistic causality always applies, but that the lattice Lieb-Robinson model is not the relevant one: their scheduling is governed by propagation, loss, detection, heralding and feed-forward. Compiled by Olivier Ezratty from the references cited in this subsection.

- **Platform-dependencies** with computing v_{eff} estimates to translate spatial distance into logical wait windows.

5.5.5 Distributed quantum computing

Distributed quantum computing (DQC) aims to scale quantum processors by networking smaller modules via quantum interconnects. While DQC can bypass some physical limitations of monolithic quantum processor architectures like chip qubit variability, frequency crowding, and crosstalk effects, it introduces new constraints related to the LR bound, which governs the speed of entanglement and information propagation in the global QPU system.

This adds to more classical light-cone constraints related to the classical control and clock coordination of qubit operations across multiple quantum processing units, which we can call a control light-cone. It is set by cable delays, FPGA and controller latency and clock distribution, all in the tens to hundreds of nanoseconds range over a few meters, and it is a classical engineering constraint rather than an LR effect. Sigal’s entanglement light cone [2362] is the quantum counterpart for the entanglement distribution step itself, giving a lower bound on the time to transport entanglement to a distant network node under ideal lossless conditions. Both point in the same architectural direction, namely a shift from global unitaries spanning multiple processors towards Local Operations and Classical Communication (LOCC) between them.

Microwave interconnect

In superconducting quantum processors, modules can be connected via microwave transmission lines which are either coaxial cables or flexible waveguides of length $L \sim 1\text{--}2\text{ m}$. These links support entanglement distribution via flying microwave photons or shared resonators, often connected to tunable couplers.

The LR bound manifests here as a constraint on the time required to establish correlations between modules:

$$t_{\text{link}} \geq \frac{L}{v_g}$$

which is the one-way propagation term only, to which device, transduction, heralding and classical control latencies must be added separately, where $v_g \approx 0.6c \approx 2 \times 10^8\text{ m/s}$ is the microwave group velocity in the cable. For $L = 2\text{ m}$, this gives a one-way propagation delay of about 10 ns, and a round trip of 20 ns, which is comparable to a single-qubit gate and small against transmon coherence times ($T_2 \sim 100\text{ }\mu\text{s}$). Applying the on-chip $v_{\text{LR}} \sim 10^5\text{ m/s}$ to a cable would give 20 μs , but that is a category error: a transmission line supports a propagating mode and is not a chain of locally coupled lattice sites, so the LR bound does not govern it. The real cost of a microwave interconnect is not propagation but insertion loss, transfer fidelity and the entanglement distribution rate, plus the classical latency of the heralding and feedforward loop, which is what actually demands careful gate synchronization in the compiler.

Optical photonic interconnect

Photonic interconnects—using optical fibers and single-photon sources—offer lower losses and longer range. In recent experiments [2395], trapped-ion modules separated by $L = 2$ m were entangled via heralded photonic Bell pairs, enabling gate teleportation with 86% fidelity. Here, the LR bound is less restrictive because:

- Photons travel at $c/n \sim 2 \times 10^8$ m/s in fiber.
- Entanglement is established before gate execution, decoupling initial entanglement propagation from computation.
- Quantum memories buffer entangled states, allowing asynchronous gate teleportation.

The effective LR velocity is replaced by the photonic link latency and memory coherence time. For quantum memories with $T_2 \sim 1$ ms, pre-generated entangled pairs can be buffered and consumed on demand. A Bell pair used for teleportation or for a remote entangling gate is consumed by that operation, so what a memory buys is the decoupling of generation from use, not the reuse of the same pair. Consumed pairs have to be replenished at a rate consistent with the workload.

The LR bound is still applicable, even when there's some pre-established entanglement between two QPUs via a photonic link. Then, when dealing with information, one may wonder about its quantity that is propagating between two QPUs holding each a very large Hilbert space with a single qubit to qubit connectivity. Isn't there some physical bottleneck there? Or is it a too realistic interpretation of what happens at the quantum level?

Microwave-to-optical transduction

Hybrid systems using microwave-to-optical transducers bridge superconducting qubits with photonic networks. Xie et al. [2396] reached percent-level efficiency with ^{171}Yb ions doped in a YVO_4 crystal at 340 ppm, exploiting an effective resonant $\chi^{(2)}$ nonlinearity of about 10^7 pm/V, with an added noise as low as 1.24 photons and without an engineered optical cavity. Other transducer families, based on electro-optic or piezo-optomechanical coupling, reach comparable or somewhat higher efficiencies but with different noise and bandwidth trade-offs.

The LR bound still applies within each interconnected processor, but inter-QPU operations are governed by photonic link dynamics and memory buffering with different issues, more related to efficiencies and fidelities than speed of entanglement propagation.

5.5.6 QEC and FTQC

Quantum error correction (QEC) and fault-tolerant quantum computing (FTQC) also interact with the LR

bound. LR bounds are mathematical statements about a specified dynamics. What they can do is help justify an approximate locality assumption for a given interacting model, which is a different claim from containing crosstalk or correlated noise, both of which come from hardware design and shielding. They refine local noise models and support assumptions used in fault-tolerant quantum error correction.

Surface codes rely on local error assumptions. LR bounds can support the approximate locality of the underlying dynamics under their own hypotheses, but they do not validate the stochastic, Markovian, leakage, correlation and measurement assumptions on which a threshold proof rests. Those require an explicit fault model of their own. Bosonic codes use cavity modes to encode quantum information. LR bounds help isolate modes, suppress photon hopping, and limit error propagation. LR bounds also guide cavity design and pulse shaping to minimize non-local interactions and correlated errors.

Error correction and locality constraints

QECC like the surface code rely on local stabilizer measurements. Each stabilizer is weight-4 and all of them are measured in parallel, so a syndrome extraction cycle is a constant-depth circuit whose duration does not grow with the code distance d . On superconducting hardware that cycle is about $1 \mu\text{s}$, dominated by readout and reset rather than by any propagation effect. Over that same microsecond an LR light cone at 10^5 m/s would sweep 10 cm, or 500 lattice sites at a $200 \mu\text{m}$ pitch, so the LR bound sits far from the binding constraint.

The code distance does play a role elsewhere. A logical operator spans d sites, so a logical operation that has to move or deform a patch of diameter d obeys:

$$t_{\text{logical}} \geq \frac{d \cdot a}{v_{\text{eff}}}$$

where v_{eff} is the circuit light-cone velocity a/t_{gate} , not the free-evolution v_{LR} . For $d = 25$ and a 200 ns gate this gives about $5 \mu\text{s}$, roughly the d rounds of syndrome extraction that lattice surgery requires anyway. The other place distance matters is classical: the decoder has to keep up with the syndrome stream, which is a classical throughput problem and the one that actually worries architects. The LR bound is what validates the locality assumption underlying the threshold theorem [2365], and that is its real role in QEC, rather than setting a cycle time.

Work by Friedman et al. in 2022 in [2393] shows that even with active feedback and measurements, the LR velocity can only be enhanced by a factor of $M + 1$, where M is the number of measured regions. Thus, even with aggressive parallelism, the speed-up is bounded:

$$v_{\text{eff}} \leq (M + 1)v_{\text{LR}}$$

Magic state distillation and gate teleportation

Topological codes offer some transversal gates, which are inherently fault-tolerant and local, but not the same ones. The surface code has a transversal CNOT between patches and a Hadamard up to a lattice rotation, while the S gate is not transversal and needs fold-transversal constructions or extra resources. Two-dimensional color codes do have the full transversal Clifford group, at the price of higher-weight stabilizers. The Eastin-Knill theorem [505] rules out any code with a universal transversal gate set, which is exactly why non-Clifford gates need a separate mechanism.

However, non-Clifford gates like T gates require additional resources, often invoking magic state distillation or cultivation and gate teleportation with their own latencies. The means of implementing these gate teleportation to implement non-Clifford gates transversally should probably account for the LR velocity.

Magic state distillation (MSD) is a protocol for preparing high-fidelity non-Clifford resource states, such as the T-state:

$$|T\rangle = \frac{|0\rangle + e^{i\pi/4}|1\rangle}{\sqrt{2}}$$

Standard magic state injection consumes an encoded magic state, applies an entangling Clifford between it and the target logical qubit, measures, and applies a measurement-dependent Clifford correction, which is often absorbed in the Pauli frame. A separate Bell pair is required only in the variants where the resource state has to be moved between distant patches first, so the exact circuit being referenced should be stated with its resource list.

Let L be the spatial separation between logical qubits. The LR bound implies that the entangled Bell pair must be distributed within time:

$$t_{\text{teleport}} \geq \frac{L}{v_{\text{LR}}}$$

For $L = 1$ cm, this gives $t_{\text{teleport}} \geq 100$ ns, which is feasible but tight in large-scale systems.

Conventional distillation has a fixed circuit schedule, so its duration is known, while its accept or reject outcome is probabilistic. Four quantities should be kept apart: the circuit duration, the acceptance probability, the retry distribution, and the buffered throughput of a factory that runs several units in parallel. Magic states distillation itself requires multiple rounds of distillation. For example, the 15-to-1 Bravyi-Kitaev protocol reduces error from p to $\mathcal{O}(p^3)$. Each round involves Clifford op-

erations and measurements across logical qubits. The LR bound constrains the time per round:

$$t_{\text{MSD}} \sim \frac{d_{\text{logical}}}{v_{\text{LR}}}$$

where d_{logical} is the diameter of the logical qubit patch.

Recent advances propose measurement-free magic state distillation using a coherent feedback network [2397], at the cost of a reduced error suppression per round of $\mathcal{O}(p^2)$ instead of $\mathcal{O}(p^3)$. The benefit is not an LR one, it is determinism: the distillation time becomes known in advance, so magic states can be produced in sync with the logical clock cycle rather than waiting for a non-deterministic post-selection round to succeed. Magic state cultivation [2398] attacks the same overhead from the other end, growing a T state directly inside a small code patch instead of distilling it.

Light-cone reasoning also shows up on the error mitigation side. D’Alberto uses the spreading of Pauli operators, what he calls the Pauli lightcone, to derive information-theoretic limits on mitigation that go beyond the usual autocorrelation-based estimates [2399]. The underlying idea is the same as in the correction case: an observable can only be affected by errors inside its causal cone, which bounds both what has to be corrected and what can be recovered by post-processing.

Ultimately, while QEC and fault tolerance do not eliminate the Lieb-Robinson bound, they provide architectural strategies to work around it, especially in the context of non-Clifford gate synthesis and logical qubit connectivity.

5.5.7 Quantum parallelism

Quantum parallelism is often invoked as a core advantage of quantum computing. By placing a quantum system in a superposition described by a very large set of complex amplitudes in the register vector state, it is possible to modify all these amplitudes in parallel with unitary transformations. From the classical computing standpoint, it looks like a single quantum gate could modify “instantaneously” 2^N complex numbers.

However, this notion is abstract and doesn’t account for physical constraints. LR bounds challenge this view by revealing a more nuanced picture where information and correlations in spatially extended quantum systems propagate at a finite speed, even under perfectly coherent unitary evolution and without any noise.

In the circuit model, gates are usually applied synchronously across the register on multiple qubits, and the global wavefunction is supposed to update instantaneously. The number of simultaneous gates applicable to a set of N qubits depends on the quantum processor implementation. In solid state qubits like superconducting qubits, gates are supposed to be parallelized arbitrarily

on multiple qubits and qubit pairs. With trapped ions, it depends on the architecture.

Quantinuum Helios's 98-qubit QPU introduced in November 2025 has eight gate zones for state preparation, measurement and single-qubit operations, of which four support parallel two-qubit gates [2389].

Microwave/RF driven trapped-ions QPU supposedly allow parallelizing gates in multiple locations, depending on various constraints on local fields control. In the case of cold atoms, things can get even more complicated when using atom shuttling in a local zone that is illuminated by a single laser for the implementation of two-qubit gates, while single qubit gates are handled in another location using local laser control with light deflectors (AODs).

In a digital circuit, gate compatibility depends on shared-qubit conflicts, commutation and control crosstalk. Overlapping causal cones do not forbid applying gates in parallel. For continuous evolution under a specified local Hamiltonian, LR bounds do say that the influence of a local perturbation spreads only progressively through the register. A unitary operation supported on region X affects distant regions Y only after a time proportional to $d(X, Y)/v_{\text{LR}}$, where v_{LR} is the LR velocity.

If correlation generation is bounded by propagation speed, then simultaneous gate applications may not yield instantaneous entanglement or interference across distant parts of the register. This has the many practical consequences we saw in this text, with quantum compilers tasked to respect temporal constraints implied by LR bounds when sequencing gates whose effects span large distances, parallel gate application being limited to regions whose LR cones do not overlap within the gate duration and with various synchronization barriers needed to be inserted between entangling layers in distributed quantum architectures, all this depending on the structure of the quantum algorithm.

What limits parallel execution on a real machine is the circuit dependency graph, the control hardware and the crosstalk between simultaneous operations, and in a distributed setting the link latency and the entanglement generation rate. A Lieb-Robinson statement enters only for a specified continuous interaction model and a specified task.

The light cone cuts both ways for quantum advantage. Bravyi, Gosset and König proved an unconditional separation between constant-depth quantum and classical circuits, and the proof turns on the light cone itself: a constant-depth circuit has a bounded causal cone, which is what makes the separation provable rather than conjectural [2400]. Wu et al. extend the argument to adaptive shallow circuits, where mid-circuit measurement and feedforward reshape the effective light cone, and connect the cone structure to the boundary of classical simulability [2401]. So a bounded light cone is not only a constraint on what a device can do, it is also part of what makes an advantage claim provable.

Quantum algorithm designers may therefore need to revisit their usual assumptions about instantaneous correlation and global entanglement. Algorithms such as Quantum Phase Estimation (QPE), Hamiltonian simulation, or quantum error correction must account for locality constraints in the hardware. This could lead to new paradigms like decomposing circuits into causally local units that interact via staged entanglement, exploiting architectural symmetries to parallelize within LR-constrained zones and introducing asynchronous scheduling and adaptive delays to manage correlation fronts.

LR bounds do not invalidate quantum parallelism, they refine it. In spatially extensive architectures, the idealized vision of instantaneous global superposition must be replaced by the reality of finite-speed entanglement and information propagation.

Quantum advantage traditionally refers to tasks solvable exponentially faster on a quantum device than on classical systems. LR bounds enforce a finite speed for the spread of information and entanglement in spatially extended systems, which implies physical limits on gate synchronization and entangling operations. We may need to investigate how that would modify the scales of accelerations provided by quantum computing.

Quantum advantage must be interpreted not only in terms of algorithmic complexity but also in terms of spatial and temporal constraints. LR bounds mean that physically bigger is not always faster. Scaling a quantum device should, in principle, expand the Hilbert space exponentially. But due to LR bounds, larger systems may suffer from increased latency in entanglement distribution, gate operations across distant regions must be serialized or staged, reducing parallelism, and synchronization overhead may grow with physical system size, potentially negating part of theoretical algorithmic speedups. In this regime, system performance becomes bandwidth-limited, not just by error rates, but by finite information flow dictated by locality.

You could decompose large quantum circuits into qubit subsystems to cope with LR constraints. But quantum speedups emerge from coherent access to global entangled states and this decomposition would limit this benefit. Subsystem partitioning may force classical-style sequencing and coordination, diluting quantum coherence benefits. And smaller systems could land in the classical simulability zone. In essence, breaking the global Hilbert space into weakly coupled regions may simplify compilation and reduce cross-talk, but risks losing algorithmic power that stems from full superposition manipulation.

As a result, future quantum algorithms design may need to evolve toward causally-aware protocols, such as exploiting speedups in localized regions with minimal global coordination, time-domain code unfolding with distributing entanglement operations over time respecting LR cones and dynamic graph restructuring with

adapting circuit layout based on information propagation bottlenecks. Quantum advantage will thus become a spatio-temporal resource, where the shape of computation in space and time must align with the constraints of physics.

5.5.8 Analog quantum computing

Analog computing platforms such as Rydberg cold atom arrays in optical lattices exhibit power-law or long-range interactions. These modify the LR bound. They have interaction scaling as $1/r^6$. LR velocity can be super-linear, allowing entanglement across ~ 200 atoms in 2D within $10 \mu\text{s}$. This corresponds to the commercial offering of industry vendors like QuEra and Pasqal.

Recent work shows that in systems with competing short- and long-range couplings, correlations can spread supersonically, meaning faster than any finite light-cone velocity [2402–2404]. This is not a violation of the LR theorem. Cavity-mediated couplings are effectively all-to-all with no spatial decay, so they fall outside the class of Hamiltonians for which an LR bound is proved. The theorem holds wherever its hypotheses hold, and these systems are built precisely so that they do not. It arises in systems with competing short-range and global-range couplings. In such systems, the dynamics following a quench show a crossover.

At short and intermediate distances, propagation is causal and light-cone-like. At larger distances, global interactions dominate, leading to distance-independent correlation growth.

This behavior has been experimentally demonstrated in systems of interacting bosonic atoms coupled to cavity light fields or exhibiting long-range dipolar interactions [2403]. The global coupling effectively allows correlations to spread across the entire system without being constrained by spatial locality.

In standard LR-bound systems, the correlation function $C(r, t)$ decays exponentially outside the light cone:

$$|C(r, t)| \leq A \exp\left(-\frac{r - v_{\text{LR}} t}{\xi}\right)$$

In supersonic regimes, this decay becomes power-law or even flat:

$$|C(r, t)| \sim t^\alpha \quad (\text{independent of } r)$$

This implies an effective Lieb-Robinson velocity $v_{\text{eff}} \rightarrow \infty$, or at least unbounded in the thermodynamic limit.

Halati et al. observed this phenomenon in both one and two dimensions [2403], even in the presence of dissipation. Their simulations showed that the density-density correlations grow with a short-time power-law scaling that does not depend on spatial separation.

Similar behavior has been reported in bosonic superfluids and long-range spin chains [2381, 2405] and also, evaluated for long disordered transverse-field Ising chains that could be implemented in analog or gate-based quantum computers [2406]. A related refinement drops the assumption of a single uniform velocity: Matsuura et al. derive an inhomogeneous LR light cone whose local velocity varies across the system, and connect it to the break time of Krylov-space dynamics [2407].

Supersonic propagation could have interesting consequences on analog quantum computing. It can drive accelerated entanglement distribution which is useful for analog quantum simulations. It can also reduce circuit depth and enable algorithms requiring global correlations to execute faster. However, these regimes are often sensitive to noise and require precise control over interaction profiles. They are more accessible in analog platforms than in gate-based quantum computers.

From a practical view, and to use an analogy with the process described in Figure 303 for gate-based QPUs, you apply a global laser drive on a set of cold atoms with a laser. All the atoms start to interact with their neighbor simultaneously. But the global propagation of these interactions to the whole system is also constrained by the LR bound.

Now, let's get the numbers for an analog quantum simulation with cold atoms:

- The Van der Waals interaction between the atoms is:

$$U = \frac{C_6(n)}{R^6}$$

where $C_6(n)$ is the van der Waals coefficient dependent on the principal quantum number n , and R is the interatomic distance.

- Experimental observations show that:

$$\frac{U}{h} \sim 0.1\text{--}10 \text{ MHz}$$

depending on the values of n and R , where h is Planck's constant.

- The typical atoms lattice spacing is $a \sim 5\text{--}10 \mu\text{m}$.
- The system coherence time is $T_2 \sim 10 \mu\text{s}$ to $100 \mu\text{s}$. It is negatively impacted by laser phase noise, parasite electric fields (Stark shift effect), atom movement and potential collisions.

The following is an order-of-magnitude toy estimate and not a measured or universal platform velocity, Lieb-Robinson prefactors depending on the Hamiltonian, the geometry, the coordination number, the interaction range and decay, the operator norms and the specific bound used. With that caveat, we define the frequency scale as $\nu = U/h$. Using the approximation

for the Lieb-Robinson velocity: $v_{\text{LR}} \sim 2\pi\nu a$, for example, with: $a = 5 \text{ }\mu\text{m}$, $\nu = 1 \text{ MHz}$, we obtain: $v_{\text{LR}} \approx 2\pi \times 10^6 \text{ Hz} \times 5 \times 10^{-6} \text{ m} \approx 30 \text{ m/s}$. This corresponds to a traversal time across approximately 100 lattice sites of $t_{\text{traverse}} \sim 10\text{--}20 \text{ }\mu\text{s}$ which at this point is compatible with the best-in-class effective coherence time. But with a larger system, we may reach this limit.

How about quantum annealing which is another form of analog quantum computing? The literature is less abundant. [2408] suggests that obtaining some quantum speedup with quantum annealing is also constrained by LR bounds. The bound could be more critical if D-Wave scaled beyond their existing 5,600 and 4,400 qubits with larger chips. It would limit the effective reach of entanglement during the annealing schedule and the annealing capability to find a global minimum of the encoded Hamiltonian.

5.5.9 Quantum thermodynamics

Lieb–Robinson (LR) bounds establish a finite speed for information and energy propagation in quantum systems. This has implications for local thermalization, entropy production, and energy transport in quantum thermodynamics. Subsystems thermalize locally before reaching global equilibrium, affecting heat and entropy dynamics in quantum processors.

Specific results derived under LR-type assumptions bound energy and entropy transport for particular dynamics and thermodynamic quantities, which is not the same as a general constraint on entropy generation. Each such statement should be cited with its own theorem and hypotheses. LR bounds complement quantum speed limits (QSL), bounding not just state evolution but also energy and information propagation.

Upadhyaya et al. push this further by making the temperature dependence explicit, showing that quantum state transfer is thermally throttled. The achievable transfer speed degrades with the energy density of the medium carrying the state [2409]. This matters for the engineering argument of this subsection, since it means v_{LR} is not a fixed hardware constant but depends on the thermal state of the substrate, and it ties the LR discussion to the energy budget question treated elsewhere in this book.

In [2371], Bruno Nachtergaele from UC Davis extended in 2011 the LR framework to irreversible quantum dynamics on lattice systems, where the evolution is governed by time-dependent generators that decay suitably in space. It proved that LR bounds still hold in this more general setting and establishes the existence of the thermodynamic limit for such dynamics. It did show


that LR bounds are robust even under non-unitary, dissipative dynamics, which are common in open quantum systems and quantum error correction. The existence of a thermodynamic limit ensures that simulation strategies can scale consistently with system size, even under irreversible evolution. It also suggests that compiler logic and scheduling can still rely on LR bounds when modeling decoherence or engineered dissipation. For quantum platforms that incorporate dissipative elements like reservoir engineering and Lindbladian control, this paper implied that LR bounds can guide the design of dissipative protocols with controlled information spread. As a result, compiler strategies can incorporate dissipative steps without losing locality guarantees. It can also help with QEC, where dissipative recovery maps must preserve locality.

5.5.10 Reality

The connection between LR bounds and quantum computing may have some consequences on quantum foundation debates. They are currently focused on the interpretation of quantum measurement and measurement correlations on entangled pairs of quantum objects. But the nature of reality in a quantum system remains elusively handled. Most quantum physicists define quantum reality as being the one that comes out of measurement, the rest is relying on mathematical representations of probabilities and vectors in Hilbert spaces. Some even discuss about the role of the observer, up to the consciousness of the observer.

What follows is a personal essay on quantum foundations rather than a consequence derived from the Lieb–Robinson bound, which says nothing about the reality of a quantum state or about observers. But what happens before measurement? What is the reality of a quantum system while being subject to unitary transformation, before any measurement, even idealized without noise? How to describe the reality in a quantum system where information propagates like a wave at a given speed, but when we cannot directly observe this propagation? To what physical reality is this information corresponding? We can only obtain some post-mortem traces of these transformations through quantum measurement at different stages of the system evolution, using the commutator presented earlier in this subsection. You can reconstruct quantum states and tomographies of quantum systems at different stages of their evolution, which provides a statistical view of their reality.

So, are quantum engineers developing quantum algorithms, compilers and optimizers crafting solutions for systems that have no palpable reality? You have 4 hours...!



5.5.11 Bibliography and notes

These are the bibliographical references and notes for the **Lieb-Robinson bound** subsection.

- [505] [Restrictions on Transversal Encoded Quantum Gate Sets](#), by Bryan Eastin and Emanuel Knill, Physical Review Letters, March 2009 (4 pages) (preprint on [arXiv](#)).
- [2226] [Roadmap on Quantum Thermodynamics](#), by Steve Campbell, Irene D'Amico, Mario A. Ciampini, Janet Anders, Natalia Ares, Alexia Auffèves, Michele Campisi, Sebastian Deffner, Cyril Elouard, Mauro Paternostro, Jukka Pekola, Janine Splettstoesser, Nicole Yunger Halpern, et al., arXiv, January 2026 (65 pages).
- [2361] [The finite group velocity of quantum spin systems](#), by Elliott H. Lieb and Derek W. Robinson, Communications in Mathematical Physics, May 1972 (9 pages).
- [2362] [Light-Cone Structure of Propagation of Entanglement](#), by Israel Michael Sigal, arXiv, April 2026 (19 pages).
- [2363] [Lieb–Robinson bounds for open quantum systems with long-ranged interactions](#), by Ryan Sweke, Jens Eisert, and Michael Kastner, Journal of Physics A: Mathematical and Theoretical, September 2019.
- [2364] [An improved Lieb-Robinson bound for many-body Hamiltonians with power-law interactions](#), by Dominic V. Else, Francisco Machado, Chetan Nayak, and Norman Y. Yao, arXiv, September 2018 (26 pages).
- [2365] [Lieb-Robinson bounds, automorphic equivalence and LPPL for long-range interacting fermions](#), by Stefan Teufel and Tom Wessel, arXiv, July 2025 (34 pages).
- [2366] [Light-cone-like spreading of correlations in a quantum many-body system](#), by Marc Cheneau, Peter Barmettler, Dario Poletti, Manuel Endres, Peter Schauß, Takeshi Fukuhara, Christian Gross, Immanuel Bloch, Corinna Kollath, and Stefan Kuhr, Nature, 2012 (preprint on [arXiv](#)).
- [2367] [Lieb-Robinson causality and non-Fermi liquids](#), by Subham Dutta Chowdhury, Sean A. Hartnoll, and Aditya Hebbar, arXiv, October 2025 (17 pages).
- [2368] [Quasiparticle engineering and entanglement propagation in a quantum many-body system](#), by Petar Jurcevic, Ben P. Lanyon, Philipp Hauke, Cornelius Hempel, Peter Zoller, Rainer Blatt, and Christian F. Roos, Nature, July 2014 (4 pages) (preprint on [arXiv](#)).
- [2369] [Non-local propagation of correlations in quantum systems with long-range interactions](#), by Philip Richerme, Zhe-Xuan Gong, Aaron Lee, Crystal Senko, Jacob Smith, Michael Foss-Feig, Spyridon Michalakis, Alexey V. Gorshkov, and Christopher Monroe, Nature, July 2014 (4 pages) (preprint on [arXiv](#)).
- [2370] [Probing quantum information propagation with out-of-time-ordered correlators](#), by Jochen Braumüller, Amir H. Karamlou, Yariv Yanay, Bharath Kannan, David Kim, Morten Kjaergaard, Alexander Melville, Bethany M. Niedzielski, Youngkyu Sung, Antti Vepsäläinen, Roni Winik, Jonilyn L. Yoder, Terry P. Orlando, Simon Gustavsson, Charles Tahan, and William D. Oliver, Nature Physics, January 2022 (7 pages) (preprint on [arXiv](#)).
- [2371] [Lieb-Robinson Bounds and Existence of the Thermodynamic Limit for a Class of Irreversible Quantum Dynamics](#), by Bruno Nachtergaele, Anna Vershynina, and Valentin A. Zagrebnov, arXiv, April 2011 (15 pages).
- [2372] [Locality and Digital Quantum Simulation of Power-Law Interactions](#), by Minh C. Tran, Andrew M. Childs, et al., PRX, July 2019 (22 pages).
- [2373] [Lieb-Robinson bounds for Bose-Hubbard Hamiltonians: A review with a simplified proof](#), by Marius Lemm and Carla Rubiliani, arXiv, March 2026 (28 pages).
- [2374] [Tightening the Lieb-Robinson Bound in Locally Interacting Systems](#), by Zhiyuan Wang and Kaden R.A. Hazzard, PRX Quantum, September 2020 (24 pages).
- [2375] [Nearly Linear Light Cones in Long-Range Interacting Quantum Systems](#), by Michael Foss-Feig, Zhe-Xuan Gong, Charles W. Clark, and Alexey V. Gorshkov, Physical Review Letters, April 2015 (5 pages) (preprint on [arXiv](#)).
- [2376] [Strictly Linear Light Cones in Long-Range Interacting Systems of Arbitrary Dimensions](#), by Tomotaka Kuwahara and Keiji Saito, Physical Review X, July 2020 (19 pages).
- [2377] [Hierarchy of Linear Light Cones with Long-Range Interactions](#), by Minh C. Tran, Chi-Fang Chen, Adam Ehrenberg, Andrew Y. Guo, Abhinav Deshpande, Yifan Hong, Zhe-Xuan Gong, Alexey V. Gorshkov, and Andrew Lucas, Physical Review X, July 2020 (19 pages).
- [2378] [Finite Speed of Quantum Scrambling with Long Range Interactions](#), by Chi-Fang Chen and Andrew Lucas, Physical Review Letters, December 2019 (6 pages).
- [2379] [Lieb-Robinson Bounds and the Generation of Correlations and Topological Quantum Order](#), by Sergey Bravyi, Matthew B. Hastings, and Frank Verstraete, Physical Review Letters, July 2006 (4 pages).
- [2380] [Quasi-Locality Bounds for Quantum Lattice Systems. Part I. Lieb-Robinson Bounds, Quasi-Local Maps, and Spectral Flow Automorphisms](#), by Bruno Nachtergaele, Robert Sims, and Amanda Young, arXiv, March 2019 (106 pages).
- [2381] [Lieb-Robinson bounds and out-of-time order correlators in a long-range spin chain](#), by Luis Colmenarez and David J. Luitz, Physical Review Research, October 2020 (9 pages).
- [2382] [Lieb-Robinson bounds and the simulation of time evolution of local observables in lattice systems](#), by Martin Kliesch, Christian Gogolin, and Jens Eisert, arXiv, August 2014 (13 pages).
- [2383] [Lieb-Robinson bounds for commutator-bounded operators](#), by Isabeau Prémont-Schwarz, Alioscia Hama, Israel Klich, and Fotini Markopoulou-Kalamara, arXiv, April 2010 (5 pages).
- [2384] [Quasi-adiabatic Continuation for Disordered Systems: Applications to Correlations, Lieb-Schultz-Mattis, and Hall Conductance](#), by M. B. Hastings, arXiv, January 2010 (33 pages).
- [2385] [Observation of constructive interference at the edge of quantum ergodicity](#), by Google Quantum AI, Collaborators, and Dmitry A. Abanin, Nature, October 2025 (76 pages).
- [2386] [Lieb-Robinson Bound and Locality for General Markovian Quantum Dynamics](#), by David Poulin, Physical Review Letters, May 2010 (4 pages) (preprint on [arXiv](#)).
- [2387] [Quantum speed limits for quantum-information-processing tasks](#), by Jeffrey M. Epstein and K. Birgitta Whaley, Physical Review A, April 2017 (11 pages) (preprint on [arXiv](#)).
- [2388] [When Does a Spin Flip? Arrival Time Distributions and Information Propagation in Discrete Quantum Systems](#), by Lionel Martellini, Entropy, March 2026 (16 pages).
- [2389] [A 98-qubit trapped-ion quantum computer with all-to-all connectivity](#), by Anthony Ransford, M. S. Allman, and Chris Langer, Nature, July 2026 (12 pages).
- [2390] [Persistence of Locality in Systems with Power-Law Interactions](#), by Zhe-Xuan Gong et al., PRL, July 2014.
- [2391] [Efficient Long-Range Entanglement Using Dynamic Circuits](#), by Elisa Bäumer, Vinay Tripathi, Derek S. Wang, Patrick Rall, Edward H. Chen, Swarnadeep Majumder,

- Alireza Seif, and Zlatko K. Mineev, PRX Quantum, August 2024 (22 pages).
- [2392] [Resilience Beyond the Light Cone: Error-Detected Primitives for Practical Dynamic Circuits](#), by Kevin C. Smith, Bibek Pokharel, Satvik Maurya, and Maika Takita, arXiv, September 2026 (26 pages).
 - [2393] [Locality and error correction in quantum dynamics with measurement](#), by Aaron J. Friedman, Chao Yin, Yifan Hong, and Andrew Lucas, arXiv, July 2022 (89 pages).
 - [2394] [Time Independence Does Not Limit Information Flow. I. The Free-Particle Case](#), by Dong Yuan, Chao Yin, T. C. Mooney, Christopher L. Baldwin, Andrew M. Childs, and Alexey V. Gorshkov, arXiv, May 2025 (22 pages).
 - [2395] [Distributed quantum computing across an optical network link](#), by D. Main, D. M. Lucas, et al., Nature, February 2025 (14 pages).
 - [2396] [Scalable microwave-to-optical transducers at the single photon level with spins](#), by Tian Xie, Rikuto Fukumori, Jiahui Li, and Andrei Faraon, Nature Physics, January 2025 (8 pages)  (preprint on [arXiv](#)).
 - [2397] [Magic state distillation without measurements and post-selection](#), by Sascha Heußen, APL Quantum, December 2025 (13 pages).
 - [2398] [Magic state cultivation: growing T states as cheap as CNOT gates](#), by Craig Gidney, Noah Shutty, and Cody Jones, arXiv, September 2024 (41 pages).
 - [2399] [The Pauli Lightcone: Information-Theoretic Error Mitigation Beyond the Autocorrelation](#), by Paolo D’Alberto, arXiv, August 2026 (19 pages).
 - [2400] [Quantum advantage with shallow circuits](#), by Sergey Bravyi, David Gosset, and Robert König, Science, October 2018 (4 pages)  (preprint on [arXiv](#)).
 - [2401] [Quantum Advantage with Adaptive Shallow Circuits](#), by Yusen Wu, Yukun Zhang, Xiaoming Zhang, Chuan Wang, and Xiao Yuan, arXiv, August 2026 (29 pages).
 - [2402] [Speed limits and locality in many-body quantum dynamics](#), by Chi-Fang Chen, Andrew Lucas, and Chao Yin, Reports on Progress in Physics, September 2023 (53 pages).
 - [2403] [From Light-Cone to Supersonic Propagation of Correlations by Competing Short- and Long-Range Couplings](#), by Catalin-Mihai Halati, Ameneh Sheikhan, Giovanna Morigi, Corinna Kollath, and Simon B. Jäger, arXiv, February 2025 (12 pages).
 - [2404] [Lieb-Robinson bounds with exponential-in-volume tails](#), by Ben T. McDonough, Chao Yin, Andrew Lucas, and Carolyn Zhang, arXiv, February 2025 (32 pages).
 - [2405] [Light-cone effect and supersonic correlations in one- and two-dimensional bosonic superfluids](#), by Giuseppe Carleo, Federico Becca, Michele Fabrizio, et al., Physical Review A, March 2014  (preprint on [arXiv](#)).
 - [2406] [The Lieb-Robinson correlation function for long disordered transverse-field Ising chains](#), by Brendan J. Mahoney and Craig S. Lent, arXiv, December 2025 (13 pages).
 - [2407] [Krylov Break Times from an Inhomogeneous Lieb-Robinson Light Cone](#), by Shunji Matsuura, Yoji Kawamura, Joseph Salfi, and Satoshi Iso, arXiv, August 2026 (23 pages).
 - [2408] [Tight Lieb-Robinson Bound for approximation ratio in Quantum Annealing](#), by Arthur Braid, Simon Martiel, and Ioan Todinca, arXiv, November 2023 (21 pages).
 - [2409] [Thermal Throttling of Quantum State Transfer](#), by Twesh Upadhyaya, T. C. Mooney, Yifan Hong, and Alexey V. Gorshkov, arXiv, August 2026 (18 pages).

This **Lieb-Robinson bound** subsection contains 51 bibliographical references and notes containing at least 1,075 pages.

5.6 Quantum uncertainty

Estimating when scalable and useful quantum computers will be available is a difficult art and science. Assessments spread between optimism and pessimism. Some entrepreneurs expect to achieve miracles in less than one decade while some scientists, on the other hand, think that this will never happen. In between, other scientists are moderate optimists and expect the wait to last at least a couple decades. We will look here at these various opinions.

Where is the chasm in these opinions? I see several sources...

Cultural bias. Americans may be natural optimists while Europeans and particularly French or Germans may be pessimistic. In some cultures, you may be pessimistic in private and optimistic in public.

Funding bias. Entrepreneurs and vendors have a tendency to oversell their offerings and roadmap to get private and public funding. That's the rule of the game in any industry. Scientists are not immune and also need to oversell to get public funding. This is twisting the "realism" game. You also hear statements according to which: "*this is the year of quantum*", "*we are seeing exponential progress*", and "*we are reaching an inflection curve*" (I wonder sometimes how to build an exponential and an inflection curve on the same chart, it would at least mean that progress stalled for many years, which is not really true). This means nothing since the related metrics are never defined. And it's false in many cases. Some vendors even talk about "*accelerating their roadmap*" when they were actually one to two years late in delivering past roadmap promises. I have been keeping track of these roadmaps for 8 years and the published versions can be compared year by year, which is the only way to evaluate such claims.

Knowledge bias. The more you know a scientific domain, the more you understand its challenges and the less you are prone to adopt some naive optimism. It's my own case in my 8+ years journey in quantum. The more I discover, the more I said to myself "*hey, that's way more complicated than I thought initially*". If you are a little curious, you discover a wealth of challenges. When you are even more curious, you look at the many potential solutions to these challenges, and in the case of quantum computers, there are many. If you have an infinite available time, you can thoroughly dissect all these solutions and then find their shortcomings. You end up navigating in a roller coaster of difficulty level assessment. If you keep a scientific eye, you can stay open to finding new solutions. But it doesn't prevent you from being reasonable.

Timing bias. When he was the cofounder of D-Wave, Georgie Rose enacted in 2003 his own equivalent of Moore's empirical law, Rose's Law, predicting a doubling every year of the number of qubits in a quantum computer, as shown in Figure 306 [2410]. Since 2007,

D-Wave delivered well on this promise, but progress has been sluggish for many figures of merit of gate-based quantum computers. While the number of qubits has steadily progressed, their quality has not followed on as fast. There is unfortunately no certain equivalent of Moore's law to assess the progress in quantum computing as shown in a 2023 preprint of mine [227]. For about two decades, the advent of useful quantum computers has always been from 5 years to a decade away. It is a quest with multiple parallel scientific and technology paths, and a moving target given classical computing is making significant progress in multiple domains (classical tensor networks, ability to simulate larger many body systems, improved deep learning techniques, more powerful GPUs).

Several papers showcase exponential progress in qubit count, fidelities and the like and extrapolate it in a simplistic manner. Beyond ignoring the scientific and engineering challenges, these plots mix quantities that are not normalized: physical qubit count, connectivity, error metrics, logical performance and actual availability are distinct axes, and a vendor roadmap is a plan rather than an observation. A usable dataset would date each point, separate physical from logical milestones, state the operational status of each system, and mark the provenance of every forecast. One good example is the typical plot exhibiting the evolution of the number of physical qubits over time. It often merges actual running systems and prototype systems not characterized nor available publicly on the cloud. It can account for 433 and 1,121 qubits corresponding to the Osprey and Condor IBM systems that were released in 2022 and 2023 but were useless and not used due to their very low two-qubit fidelities. As a result, they were quickly dismantled. These systems were actually test-beds for scalable control electronics and cabling, where they did fare well. Considering "exponential growth", Google and IBM had 53 qubits in 2019. In 2025, they're at respectively 105 and 120 for their best chips with 1-to-4 qubit connectivity (IBM has 156 qubits with Heron r3 but with a lower qubit connectivity). In 2020, Google didn't expect that it would take them 5 years to double the number of their physical qubits. Trapped ions have good fidelities but IonQ has been stuck at 36 qubits for a long time and delivered 64 qubits in September 2025, Quantinuum is now at 98 and AQT at 20. This is not really exponential. I looked at this in a paper I published on quantum computing supposed Moore's law in 2023 [227]. You can see where we are with the relation between physical qubit number and two-qubit gate fidelity in Figure 252.

You can also understand why you need to be cautious when interpreting these "exponential charts" with looking at a similar chart created in 2015 that positioned NMR qubits as best-in-class for their scalability potential [2412] (Figure 307). In reality, NMR qubits didn't really scale at all. Some exponential law can however be observed in the evolution of other operating parameters of quantum computers such as the stability time of

Rose law (2003) "quantum Moore's law"




Figure 306: Compilation of the putative equivalents of Moore's law in quantum computing. They all need updates! Otherwise, you can't prove there's a real ongoing acceleration of progress in quantum computing science and technology. Sources: [2411] for the lower-right chart. The first one on the left is from Michel Devoret.

qubits, their error rate and the number of consecutive operations performed reliably [1395, 2413]. More conservative timelines exist, like the one from Q-CTRL which places FTQC quantum advantage around 2033 and large scale FTQC around 2040 (Figure 308). This remains a vendor forecast rather than an empirical extrapolation, and like the others it should be compared, when the date comes, against what was actually delivered.

Science vs engineering bias. Some scientists, including many Nobel laureates I had a chance to meet, consider that the science is settled and that building scalable and fault-tolerant quantum computers is just a matter of engineering. It's more subtle than this. Of course, many engineering and technology developments are needed to create FTQC. But science is still at play in a back-and-forth scheme between science and engineering. We push the limits of science with engineering to create large scale quantum computers. We are discovering ways to control even larger numbers of entangled quantum objects, without having a clear idea of the real limits we will discover. What new sources of noise will we discover (correlated, crosstalk, etc.)? Will we be able to mitigate it? How far can we reduce noise and error with qubit operations? In the algorithm landscape, the theory is still in the making to more clearly describe the conditions to obtain some theoretical and practical exponential speedups. And on and on. What may happen is the creation of some pressure on scientists to better understand the phenomena that will be uncovered thanks to large scale engineering.

As Christophe Jurczak from Quantonation argued in a 2026 essay, the constraints to build quantum computers are endogenous and technical, not linked to some sort

of market failure or even investment limitations [2415]. I would reinforce this point in saying that the scientific and technical difficulty is generally underestimated, particularly given many components like quantum error correction are being designed in parallel with the hardware, but with very optimistic assumptions (see 5.2.11).

5.6.1 Optimism

Google said it achieved quantum supremacy in October 2019. It was not solving a specific problem since it was based on a randomized benchmark with no input data. It was found later that it could be efficiently simulated classically on a classical computer, faster and with less energy thanks to the use of Nvidia GPUs and tensor networks. In a word, classical computing progress was faster than quantum computing progress. The same happened with the boson sampling experiments from China between 2019 and 2021. These were unprogrammable random photon mixers. Later boson sampling experiments in 2021 and 2022, like with Xanadu, were programmable, but did not show a practical computing advantage.

Then, starting in 2023, IBM introduced its "quantum utility" mantra based on their 127 and 133 qubit QPUs and the use of quantum error mitigation. Their many case studies were, however, quite theoretical and remote from industry needs and requirements.

As published in their roadmaps, Google, IBM and Amazon expect to achieve some true quantum advantage relatively quickly. Several quantum computing star-




Figure 307: Chart with number of qubits per technology and year, as of 2015. It gave the impression, back then, that NMR qubits were the most scalable. They are not! Source: [2412].




Figure 308: High-level FTQC roadmap according to Q-CTRL, published in 2024. Source: [2414].

tups like PsiQuantum, QuEra, Pasqal, Infleqtion and Alice&Bob expect to deliver 100 logical qubits between 2028 and 2033. A few are even mentioning the 2,000 logical qubits milestone by 2033 or so.

Kenneth Regan thought in 2017 that an industry vendor - probably Google - would claim to have reached quantum supremacy in 2018 and that it would quickly be contradicted by the scientific community [2416]. This

happened in 2019, and the contradictions came fast. That was quite a good prediction!

Alain Aspect does not see any strong scientific or theoretical obstacle preventing the creation of reliable quantum computers. Many quantum scientists, including those I met at the Lindau Nobel Laureate meetings in July 2024, believe that uncertainty is mostly a technological and engineering one. They think that the process can be accelerated, if it is fueled by good talent and public/private investments.

John Preskill has the same opinion: it will work, but it will take several decades. He went from devising the advent of the NISQ era in 2017 to the Megaqop era in 2024.

Nicolas Gisin estimated in 2022 that the pace to quantum usability was accelerating [2417].

Jian-Wei Pan was more optimistic in a 2022 interview, expecting a regular quantum advantage before 2027 [2418]. That position is relayed through a secondary source, it is four years old at the time of writing, and it does not specify which kind of advantage is meant, computational, utility-scale, economic or fault-tolerant.

Optimists also include the many hardware quantum computing startups, whose solutions are expected to work at large scale within the next five years, a horizon that has kept moving forward with each passing year. These successive roadmap versions are archived and can be compared with what was delivered.

Among the most optimistic public voices is **Michio Kaku**, a physicist turned futurist and best-selling author who cocreated string field theory. In his 2023 book “Quantum Supremacy”, he wrote that quantum computing will soon “*solve some of humanity’s biggest problems, like global warming, world hunger, and incurable disease*”. The book does not provide the resource estimates that would support such claims, and the applications it describes correspond to workloads whose logical qubit counts and non-Clifford gate counts are, on current estimates, far beyond the machines expected within a reasonable horizon.

Finally, in the optimism side, you can count on the industry vendors and analysts who say that industry-grade quantum advantage is around the corner if not already there, that it’s urgent to adopt quantum computing, and the like. The headlines are all over the place about this stance. They convey misleading messages like the upcoming capability of quantum computing to run LLMs with less energy than classical Nvidia GPU-driven data centers, or to solve traffic jams in real time. For many reasons explained in this book, this is not plausible based on our existing knowledge of the matter: quantum computers may/will be adapted to low-data/high-complexity problems, and when they reach a quantum advantage in computing time, particularly to solve combinatorial optimization problems, it will be long computing times. The more reasonable statements are those sticking to quantum chemistry and material

design applications. They have their own challenges but they are still more plausible.

5.6.2 Pessimism

Pessimism comes from a few researchers. Their doubts lie mainly in the ability to fix the noise that affects qubits, whatever its type [2419].

One of these is the Israeli researcher **Gil Kalai** who believes that we will never be able to create quantum computers with a low error rate [2420, 2421]. According to him, we cannot create stable quantum computers because of the noise that affects the qubits and “*the ultimate triumph of quantum information theory will be in explaining why quantum computers cannot be built*” [2422, 2423]. This is illustrated in the scale below in Figure 309, which sets the lowest reasonably achievable noise level well above the level required to create a scalable quantum computer.




Figure 309: Gil Kalai’s quantum computing errors complexity scale.

He has been working for a while on the creation of some mathematical model that would prove the impossibility of overriding these errors, even with quantum error correction codes. In 2022, he published another paper to prove his point, with a philosophical approach related to the notion of free will [2424].

Another skeptic of quantum computing is **Mikhail Dyakonov** (born in 1940 in the USSR) who works in the Charles Coulomb Laboratory (L2C) of the CNRS and the University of Montpellier in France. He detailed his views in an article at the end of 2018 [2425], which he later turned into a short book [2426]. His argument is more intuitive but less documented than the work of Gil Kalai and it generated some rebuttals by Ben Crige and Scott Aaronson [2427–2429].

Serge Haroche believes universal quantum computing is an unreachable dream, also because of that damned noise. He stated that claim in a famous 1996 paper coauthored by Jean-Michel Raymond [156]. The irony of this paper is it was published during a year with much progress in quantum computing Figure 56. Rainer Blatt, Juan Cirac and Peter Zoller created the first system supporting two trapped ions qubits and a CNOT gate between these qubits. Lov Grover created his fa-

mous algorithm. Peter Shor laid the basis of quantum error correction and fault-tolerant quantum computing. David DiVincenzo established his famous 5-criteria checklist to create useful quantum computers. In a way, Scott Aaronson's famous 2019 quote became quite lively. On the other hand, Serge Haroche thought and still believes that the path of analog quantum simulation, especially based on cold atoms, is reasonable and realistic.⁷

"For me, the most important application of a quantum computer is disproving the people who said it's impossible. The rest is just icing on the cake."

Scott Aaronson, 2019.

Cristian Calude and **Alastair Abbott** argue that, for the specific algorithms and oracle model they analyze, a quadratic acceleration is what remains in practice, and that classical heuristics can often match it on the instances that matter [2430]. This does not extend to quantum algorithms in general: factoring, Hamiltonian simulation, query algorithms and structured linear algebra rest on different complexity claims and different input and output assumptions, and a heuristic that works well on some instances is not a general complexity equivalence.

Liam P. McGuinness from Ulm University in Germany argues that the precision targeted by some energy estimation algorithms runs against the Heisenberg limit and the Cramér-Rao bound of quantum metrology [2431]. These bounds constrain particular estimation tasks under stated models, and turning them into an impossibility result for quantum computing would require a theorem that accounts for error correction, calibration and the chosen measurement strategy.

Xavier Waintal (CEA-IRIG in Grenoble, France) showcases serious doubts on the possibility of creating large-scale quantum computers. His culprit is noise and scale issues. Qubit operations rely on very complex n -body quantum problems. The objection that codes handle only flip and phase errors needs a qualification: for qubits, X , Y and Z form an operator basis, so syndrome extraction discretizes an arbitrary error inside the computational subspace into Pauli components, and correcting the right Pauli set corrects arbitrary in-subspace errors under the code and noise assumptions. What genuinely falls outside that picture is leakage, loss, correlated faults and non-Markovian effects, which need their own mechanisms, several of which were covered earlier. He recommends exploiting the mean-field theory which models the complex interactions between qubits and their environment [1825]. He published with others various papers that are mentioned elsewhere in this book showing the limitations of Grover, (NISQ) VQE and (FTQC) QPE algorithms. In January 2024, he published his "quantum house of cards" pamphlet summarizing his views [2432]. His arguments are both the

most documented I have seen but which, well used, may fuel interesting research to find solutions. He also exemplifies the huge progress made with classical algorithms and architectures to solve the problems envisioned for quantum computers, mainly thanks to the advances with tensor networks. Xavier Waintal uses the scale in Figure 310, with 5 difficulty levels, to build a quantum computing machine. It positions where we are right now and the challenges ahead. It goes beyond large scale computing given some quantum memory would be mandatory for some key algorithms like QML and HHL.

Amit Agar from Indiana University is describing the NISQ trap, showing how all supremacy and advantage experiments are twisting computing systems to their benefit, and are then quickly classically simulated [2433].

Leonid Levin (a Russian scientist who defined the NP complete complexity class in 1973) and **Oded Goldreich** (an Israeli professor in computer science from the Weizmann Institute) are other quantum computing skeptics mentioned by Scott Aaronson [2429]. Seth Lloyd expects FTQC QPUs to be ready by 20 years " $\pm$ infinity". Another argument against scalable quantum computing raises by Leonid Levin deals with the computational state vector amplitudes values becoming tiny as the number of qubits grows. After just applying a set of H gates on N qubits, this amplitude becomes $1/2^N$ for each computational basis state in the qubits register. It becomes quite small as N grows beyond 50. Are these values corresponding to some physical observable that would have a value way below the physical error rate in the system? Or even below some physical Planck constant? Well, it seems not. The physical observables in the system remain based on the individual qubits basis states $|0\rangle$ and $|1\rangle$, not the state vector register amplitudes. This is also discussed in [2429], but with rather unconvincing arguments on both sides of the debate.

Another proposed constraint concerns QPU clock accuracy. Imperfect reference frames introduce dephasing errors in quantum gates, and the cited works derive a tradeoff between clock resolution and accuracy under their own model [2434, 2435]. Whether this bounds scaling in practice depends on the clock architecture and on how these errors interact with error correction, so it is a modeled risk rather than a demonstrated limit.

Tim Palmer is advocating that the Hilbert space is discretized which leads him to assert that qubit information handling capacity is limited and finite, based on an idea from John Wheeler coined in 1989 [2436] and also, the Heisenberg uncertainty principle. He built a model in 2025 stating, that as a result, qubits register may not operate beyond about 600 physical or logical qubits [2437, 2438]. He landed on the conclusion that factorizing an RSA-2048 bit key would be impossible with a quantum computer. At least, this would bring some comfort to cybersecurity specialists. This is still based on various quantum mechanics conjectures, in-

difficulty scale	technology	use cases	examples
1	quantum simulator (analog-no gates)	quantum simulations	D-Wave, Pasqal
2	gates-based analog systems, low fidelity	system validation NISQ algorithms	Google, IBM, Rigetti, ..
3	gates-based analog systems, low fidelity	variational calculations in quantum chemistry	Possibly PsiQuantum
4	ideal quasi-deterministic gates-based systems (FTQC/LSQ)	factoring large numbers, exact quantum chemistry	TBD
5	4 + quantum memory	quantum machine learning, linear algebra (HHL)	TBD

Figure 310: Scale of difficulty on the way to usable quantum computers. Source: Xavier Waintal. 2021.

cluding the elusive gravitational quantum mechanics. Tim Palmer puts this under the umbrella of Rational Mechanics or RaQM [2439, 2440].

Gaspere Galati is an Italian scientist from Tor Vergata University, Rome, who has been very vocal against the quantum radar hype and fall, and made parallels with the current quantum computing hype [2441]. He advocates for improved evaluation processes of ongoing research projects.

Sabine Hossenfelder is a German physicist turned into a famous Youtuber expression blunt opinions on how research operates and skepticism on various technology developments related to physics. She is very skeptic on the advent of useful quantum computers [2442].

5.6.3 Middle road

Pessimism and optimism are moods. Let's fight the knowledge bias and overcome the fundamental research vs engineering debate, and adopt a scientific mindset.

What is difficult? What takes time? What can we discover that make things easier or more difficult in the quest for scalable quantum computers?

Here are some insights on how to put things in perspective:

- **Understand** the state of the art and the progress from the last 10 to 20 years. Avoid the abusive regression bias mentioned above. One good example of learning curve is with quantum error correction. As seen in the related part of the book, it is not an easy one. You can also track the progress in quantum algorithms improvements, particularly with the abundant work on gate-count reduction for critical tasks [1958, 2443–2445].
- **Holistically** approach the topic, covering all the stack, from qubit physics and fidelities, error correction overhead, qubit control scaling, and above all,

quantum algorithms and their benefit. In the case of superconducting qubits, look at the dilemma between reducing qubit variability across large chips (Google plans for 10K qubit chips and Qolab for 20K), scaling cabling which is not easy beyond 1K physical qubits, room temperature and cryogenic control electronics, thermal load, and signals multiplexing. Look at all the various solutions imagined by academics and industry vendors to solve the most pressing challenges. This mandates a multidisciplinary approach with mathematics, fundamental quantum physics, thermodynamics and chemistry, algorithms and software engineering. You may not do that alone and need to find the right people. Figure 311 goes into some details with laying out some of these challenges, most of which being covered extensively in various parts of this document. This book shows that every tough challenge is already addressed by many researchers and usually with many competing approaches and low TRLs (technology readiness level). In such cases, I frequently make an inventory of these variations in the book, but without the ability to rank them properly and assess their viability.

- **Learn** the required figures of merit to build FTQC QPUs and how they may relax over time, for example, thanks to improved error correction codes. One good example in the case of PsiQuantum and other photonic FTQC vendors who have to fight photon loss all over their life path (generation, circuits, routing, measurement).
- **Identify** the new challenges and their upcoming constraints, like the need for QPU interconnect and their related needed figures of merit and overhead, particularly in circuit partitioning. Also, grasp the algorithmic challenges to create exponential speedups, as presented by John Preskill and Jens Eisert in a preprint published in October 2025 [2005].
- **Compare** quantum computing with classical computing in an honest manner. Understand how a




Figure 311: Full stack software and hardware challenges for QPU scalability. The color split between scientific and engineering challenges is illustrative. Many items span both categories, architecture, algorithms, cryogenics and control among them. Blue text flags what is judged here to be dominated by open scientific questions and black text what is judged to be dominated by engineering and technology work. © Olivier Ezratty, 2023-September 2024.

quantum computer and its quantum algorithms could exceed the performance of systems. Learn about the capabilities of tensor networks and existing classical optimization solvers. Understand that most existing quantum algorithms and use cases are not at all in quantum advantage regime. Understand the weaknesses of those papers publishing some algorithm running with <30 qubits with no indication of the resources needed to scale it in a quantum advantage regime [2446].

- **Differentiate** theoretical work from practical engineering advances. The disconnect between both is sometimes staggering. This is particularly true for appreciating quantum algorithms speedups. Not all quantum algorithms showcase an exponential theoretical speedup. Grover's algorithm provides a quadratic speedup in query complexity, and whether that survives end to end depends on the oracle cost, the constants and prefactors, the available parallelism, the fault-tolerance overhead and the classical baseline, which in the regimes studied so far remove the advantage. Non-exponential speedups therefore require a full end-to-end resource analysis rather than a blanket verdict: some can be valuable at sufficient scale, while oracle and error correction costs eliminate the advantage in specific regimes. The trick of the trade is that all speedups are theoretical but not yet practical. Another way to look at this is to envision, even with moderate algorithms speedups,

an energetic advantage for quantum computing, as discussed in the related part that we just saw, but it is still uncertain and will probably not happen with gate-based NISQ QPUs. These speedups are rarely documented with taking into account all the quantum computing food chain: data preparation, oracle operations, quantum memory access when it is required, quantum error correction, non-Clifford group quantum gates generation (particularly for all algorithms using a quantum Fourier transform, and there are many) and the number of shots/runs required (with or without quantum error corrections).

- **Appreciate** the time it takes to improve the technology. Most of qubit types rely on improving materials and manufacturing techniques for various kinds of chips. The design-manufacturing-testing-characterizing-testing cycle is long, lasting between 6 and 18 months depending on the technology. Components characterization qualitatively tests and selects the manufactured components. Then, the components must be assembled in their cryostats with bulky cabling, which usually takes a lot of time. This can last up to a month and in the best case down to a week. Then, to carry out the tests, the thermalization of the computer takes about 24 hours and changing the test chip takes at least 3 to 7 hours as we will discover in the section on cryogenics. There are many iterations there, which explains for example why Google was relatively late with its Willow chip.

Sometimes, the manufacturing process doesn't exist yet like for manufacturing III-V chips to host millions of the Majorana-fermion based qubits Microsoft is planning to build in "the next few years". This design, manufacture, test and learn cycle can still be optimized using various techniques like using full-stack EDAs (electronic design automation software tools) to fasten the design and digital-twin test cycle, creating dedicated clean rooms with optimized tooling like clusters (which automatically handle wafers during the manufacturing process), cryogenic chip testing, and mixing several chip variations on the same wafers and even in the same test cryostat.

- **Broaden** your technology scope. In quantum computing, it means considering seriously analog quantum simulators. Some algorithms are being evaluated to run on these systems, like the ones from Pasqal and QuEra. Quantum annealing and photonic coherent Ising machines could also bring their share of hope. The debate is still out to assess what is the quantum advantage of D-Wave with its latest annealer generation, relying on their 4,400 Zephyr-topology chip qubits.
- **Adopt** a pipeline approach if you are a potential end-user customer. Quantum computing won't be a "*one size fits all needs*" solution. It will be combined with classical methods and hardware in multiple ways. Also, there is no such thing as a miracle quantum

algorithm. That was a key learning from the Quantum Benchmarking Initiative launched by DARPA which started to consolidate resource estimates to solve interesting problems, particularly in chemistry simulations. All of these solutions combined multiple quantum algorithms, all fed with classical pre-processing tasks [2447].

- **Imagine** other new challenges at the physical level, keeping three evidence levels apart. Observed engineering limitations come first, with measured crosstalk, fabrication variability, correlated noise and control limits at scale. Modeled risks come next, such as the constraints that a Lieb-Robinson type bound places on a specified interacting model. Speculative research questions come last, such as a possible contribution of dark matter induced events to qubit error budgets, which remains a research topic without an established weight in any error model [2448].

As a conclusion of this part, we can still count on two things to reach quantum computing scalability and practicality. One is the great diversity of paths chosen by academic scientists and industry vendors engineers. This creates a sort of fault-tolerance for innovation. The second, more generally, is we can believe and bet on researchers and engineers' creativity to solve these highly complicated problems. But it is still a very generic commonplace reasoning that has not much value per se.



5.6.4 Bibliography and notes

These are the bibliographical references and notes for the **Quantum uncertainty** subsection.

- [156] [Quantum computing: dream or nightmare?](#), by Serge Haroche et al., Physics Today, 1996 (2 pages).
- [227] [Is there a Moore's law for quantum computing?](#), by Olivier Ezratty, arXiv, March 2023 (32 pages).
- [1395] [Awesome Quantum Computing Experiments: Benchmarking Experimental Progress Towards Fault-Tolerant Quantum Computation](#), by François-Marie Le Régent, arXiv, July 2025 (29 pages).
- [1825] [What determines the ultimate precision of a quantum computer?](#), by Xavier Waintal, arXiv, April 2019 (6 pages).
- [1958] [How to factor 2048 bit RSA integers with less than a million noisy qubits](#), by Craig Gidney, arXiv, May 2025 (40 pages).
- [2005] [Mind the gaps: The fraught road to quantum advantage](#), by Jens Eisert and John Preskill, arXiv, October 2025 (11 pages).
- [2410] [Quantum computing Rose's Law is Moore's Law on steroids](#), by Matthew Griffin, Fanatical Futurist, August 2016.
- [2411] [Technical Roadmap for Fault-Tolerant Quantum Computing](#), by Amir Fruchtmann and Iris Choi, University of Oxford, October 2016 (30 pages).
- [2412] [Recent advances in nuclear magnetic resonance quantum information processing](#), by Ben Criger, Gina Passante, Daniel Park, and Raymond Laflamme, The Royal Society Publishing, 2015 (16 pages) 📄.
- [2413] [Forecasting timelines of quantum computing](#), by Jaime Sevilla and C. Jess Riedel, arXiv, December 2020 (23 pages).
- [2414] [Q-CTRL, 2024 year in review](#), by Michael J Biercuk, December 2024.
- [2415] [The Five-Year Horizon: Quantum Computing and the Philosophy of Technical Becoming](#), by Christophe Jurczak, Philosophies, July 2026.
- [2416] [Predictions we didn't make](#), by Richard Lipton, January 2018.
- [2417] [Quantum computing at the quantum advantage threshold: a down-to-business review](#), by A.K. Fedorov, N. Gisin, S.M. Belousov, and A.I. Lvovsky, arXiv, March 2022 (55 pages).
- [2418] [Jian-Wei Pan Sees Routine Quantum Advantage Within Five Years](#), by Matt Swayne, The Quantum Insider, February 2022.
- [2419] [The different forms of quantum computing skepticism](#), by Boaz Barak, 2017.
- [2420] [Why Quantum Computers Cannot Work](#), by Gil Kalai, October 2013 (60 slides).
- [2421] [How Quantum Computers Fail: Quantum Codes, Correlations in Physical Systems, and Noise Accumulation](#), by Gil Kalai, arXiv, June 2011 (16 pages).
- [2422] [The Argument Against Quantum Computers](#), by Katia Moskwitch, February 2017.

- [2423] [Further Statistical Study of NISQ Experiments](#), by Gil Kalai, Tomer Shoham, and Carsten Voelkmann, arXiv, December 2025 (30 pages).
- [2424] [Quantum Computers, Predictability, and Free Will](#), by Gil Kalai, arXiv, April 2022 (32 pages).
- [2425] [The Case Against Quantum Computing](#), by Mikhail Dyakonov, IEEE Spectrum, November 2018.
- [2426] [Will We Ever Have a Quantum Computer?](#), by Mikhail Dyakonov, Springer, 2020 📄.
- [2427] [The Case Against 'The Case Against Quantum Computing'](#), by Ben Crige, January 2019.
- [2428] [Happy New Year! My response to M. I. Dyakonov](#), by Scott Aaronson, Scott Aaronson blog, 2013.
- [2429] [Skepticism of Computing](#), by Scott Aaronson and Chris Granade, 2008.
- [2430] [The development of a scientific field](#), by Alastair Abbott et al., June 2016.
- [2431] [The Quantum Cramér-Rao lower bound \(Why quantum computers won't work I\)](#), by Liam P. McGuinness, arXiv, March 2025 (19 pages).
- [2432] [The quantum house of cards](#), by Xavier Waintal, PNAS, December 2023 (7 pages) 📄 (preprint on [arXiv](#)).
- [2433] [The NISQ Trap: Eight Years of Demonstrations the Hardware Was Built to Lose](#), by Amit Hagar, arXiv, July 2026 (8 pages).
- [2434] [Fundamental accuracy-resolution trade-off for timekeeping devices](#), by Florian Meier, Emanuel Schwarzahns, Paul Erker, and Marcus Huber, arXiv, November 2023 (7 pages).
- [2435] [Impact of Imperfect Timekeeping on Quantum Control](#), by Jake Xuereb, Marcus Huber, et al., PRL, October 2023 (13 pages) 📄 (preprint on [arXiv](#)).
- [2436] [Information, physics, quantum: the search for links](#), by John Wheeler, 3rd Int. Symp. Foundations of Quantum Mechanics, 1989 (15 pages).
- [2437] [Testing Quantum Mechanics with Quantum Computers: Qubit Information Capacity](#), by Tim Palmer, arXiv, September 2025 (11 pages).
- [2438] [Rational quantum mechanics: Testing quantum theory with quantum computers](#), by Tim Palmer, Proceedings of the National Academy of Sciences, March 2026.
- [2439] [Impossible Counterfactuals, Discrete Hilbert Space and Bell's Theorem](#), by Tim Palmer, arXiv, January 2026.
- [2440] [Solving the Mysteries of Quantum Mechanics: Why Nature Abhors a Continuum](#), by Tim Palmer, arXiv, February 2026.
- [2441] [Lesson learnt from the rise and fall of quantum radar research](#), by Gapare Galati et al., Academia Quantum, March 2025.
- [2442] [In which I lose faith in quantum computing](#), by Sabine Hossenfelder, Youtube, September 2025.
- [2443] [Block encoding with low gate count for second-quantized Hamiltonians](#), by Diyi Liu, Chao Yang, et al., arXiv, October 2025 (32 pages).
- [2444] [Accelerating Grover Adaptive Search: Qubit and Gate Count Reduction Strategies with Higher-Order Formulations](#), by Yuki Sano, Kosuke Mitarai, Naoki Yamamoto, and Naoki Ishikawa, arXiv, May 2024 (11 pages).
- [2445] [Optimal Hadamard gate count for Clifford+T synthesis of Pauli rotations sequences](#), by Vivien Vandaele, Simon Martiel, Simon Perdrix, and Christophe Vuillot, arXiv, February 2024 (32 pages).
- [2446] [How to analyze quantum computing use cases](#), by Olivier Ezratty, France Quantum, June 2025 (31 minutes).
- [2447] [Quantum Benchmarking Initiative \(QBI\) Proposers Day](#), by DARPA, DARPA, August 2024.
- [2448] [Quantum Computers as Dark Matter Detectors](#), by Aaron S. Chou, Fermilab, August 2022 (57 slides).

This **Quantum uncertainty** subsection contains 45 bibliographical references and notes containing at least 508 pages.

5.7 Economics

We already dealt a little with the economics of quantum computing in the energetics part, considering that the power consumption was the key part of the operating expenses of running quantum computers and that a key economic metric was the ratio between the economic benefit and the total cost of ownership, which is a benefit-cost ratio rather than a return on investment in the strict accounting sense (Figure 299). Contrary to classical computing technologies which are actually deployed at scale, it is difficult here to assess both the investment and the value generated by quantum computers.

We are at the very early stage of the quantum computing era. The current NISQ market is too small to generate economies of scale that are usually fueling the growth of innovative technology sectors. We can still make some projections based on a couple assumptions.

You may wonder why this economics part sits in the “engineering” section of the book. One reason is that engineering usually accounts for economic constraints for the buildout of any technology. The second is that we look here at the cost structure of quantum computers at the technical level.

5.7.1 Pricing

Quantum computer costs will be accounted for differently depending on the user organization. It may be a combination of capex (capital expenses) and opex (operating expenses) when the hardware is acquired, which happens for governments, defense and some industry verticals.

Otherwise, when the resource is consumed through the cloud on a subscription or per hour basis, the cost is ordinarily an operating expense (opex) for the customer and not a capital expenditure (capex), an owned system being the case that generates capex, subject to the buyer’s accounting policy, plus its own opex (Table 15). There are of course other indirect immaterial costs to consider like training, software development and acquisition.

Let’s first consider the hardware cost. Numbers are starting to get public, particularly as some quantum computers are acquired in the public sector through public RFPs (request for proposals), particularly as part of the EuroHPC Joint Undertaking.

The figures mentioned below are of different natures and are not necessarily comparable. Some are contract values awarded through public procurement, others are cumulative partnership values covering several system generations, purchase orders, bills of materials, total cost of ownership, or author estimates. The nature of each is stated inline.

- **D-Wave.** For a while, the only officially “priced” quantum computers on the market today were those from D-Wave. Their unit price sits between \$12M and \$20M depending on the generation and on the bundled services: the Advantage sold to the Jülich Supercomputing Centre generated \$12.6M of revenue in Q1 2025, a system was contracted in Italy for €10M, and Florida Atlantic University signed a \$20M agreement in January 2026 for an Advantage2 bundled with a training academy. They sold only a few of these. Most D-Wave customers are using D-Wave computers on the cloud in D-Wave premises. Some customers pay over \$200K per year to benefit from premium access to these machines.
- **IBM** is selling its QCs in full-stack packages of about \$40M to \$50M, even when the “sold” system is installed in their own facilities, like in Ehningen, Germany. IBM sold and/or installed a couple QCs elsewhere in South Korea, Japan, Canada and Spain, frequently in partner’s locations, various local research, university and industry partners. On top of this, they now run about a dozen QCs installed in their Poughkeepsie datacenter in New York State. For a while, they had up to 24 QCs running there but reduced their numbers when they increased their qubit numbers above 127 and, above all, drastically improved their efficiency to run quantum circuits. IBM started to sell and install actual QCs in third-party sites in 2023, like in Cleveland Clinic, Ohio.
- **IonQ** sold two systems to QuantumBasel in Switzerland, contracted in June 2023, an #AQ 35 machine then a Forte Enterprise, with a Tempo added in December 2025, bringing the cumulative value of that partnership above \$60M for four system generations, which is a partnership total spread over several years and multiple quantum processors. Their largest documented deal remains the \$54.5M contract signed in 2024 with the US Air Force Research Lab. They expect that by the end of the decade, they will build a 2 million physical qubits system delivering 40K to 80K logical qubits with a bill of materials below \$30M in 2025 dollars. This is optimistic from the feasibility standpoint, and a bill of materials (BOM) is not a price. The company will then apply a markup to the BOM, which can’t be estimated at this point.
- **AQT** delivered its 20-qubit trapped ions QC to Poland’s PSNC in 2025 for €12.28M [2470], under a EuroQCS-Poland contract signed in July 2024 and funded half by the EuroHPC JU and half by Poland. That amount covers the delivery and its associated services, while the total cost of ownership announced was €15.5M.
- **Pasqal** sold a 140-qubit analog quantum simulator to CINECA in Italy in 2025 as part of the EuroHPC JU for €13M, about \$14M [2471], with an upgrade to a hybrid analog/digital mode planned for 2027.

Vendor	Cloud service	Quantum processor	Qubits #	Qubit type and model	Price (\$/QPU-h)	Pricing conditions	Available since
 Google	Google Quantum Computing Service	Willow	105	Superconducting transmon, gate-based	Not public	Restricted to approved academic and industry users through the 2026 Early Access program which is not a public cloud access.	2026, early access
 IBM	IBM Quantum	Heron r3	156	Superconducting transmon, gate-based	2,880–5,760 ^b	Billed by QPU execution time. Pay-as-you-go: \$96/min, 1-s minimum. Flex: \$72/min with at least 400 min/year. Premium: \$48/min with at least 5,200 min/year.	July 2025
 D-Wave	D-Wave Leap	Advantage2	4,400+	Superconducting flux qubits, quantum annealing	Not public	Customer plan. QPU access time against an account quota.	May 2025
 IONQ	AWS Braket	IonQ Forte Enterprise 1	36	Trapped-ion, gate-based	7,000	Braket Direct dedicated reservation, 1-h increments. On demand: \$0.30/task + \$0.080/shot.	2025-03-20
 QuEra	AWS Braket	QuEra Aquila	256	Neutral Rb atoms, analog Hamiltonian simulation	2,500	Braket Direct dedicated reservation, 1-h increments. On demand: \$0.30/task + \$0.010/shot.	2022-11-01
 QAT	AWS Braket	AQT IBEX-Q1	12	Trapped-ion, gate-based	4,800	Braket Direct dedicated reservation, 1-h increments. On demand: \$0.30/task + \$0.0235/shot.	2025-11-20
 IQM	AWS Braket	IQM Emerald	54	Superconducting transmon, gate-based	4,000	Braket Direct dedicated reservation, 1-h increments. On demand: \$0.30/task + \$0.00160/shot.	July 2025
 rigetti	AWS Braket	Rigetti Cepheus-1-108Q	108	Superconducting, tunable couplers, gate-based	4,100	Braket Direct dedicated reservation, 1-h increments. On demand: \$0.30/task + \$0.000425/shot.	April 2026
 Pasqal	Azure Quantum	Pasqal FRESNEL, Orion Beta	100	Neutral Rb atoms, analog	3,481 ^a	Pay as you go at EUR 3,000/QPU-h plus Azure infrastructure costs. This is QPU execution time, not an exclusive one-hour reservation.	June 2023 or earlier
 Quantinuum	Azure Quantum	Quantinuum H2-2	56	Trapped ¹⁷¹ Yb ⁺ ions, QCCD, gate-based	n/a	Billed in Hardware Quantum Credits (HQC). Standard: \$125k/month for 10k H2 HQCs. Premium: \$175k/month for 17k H2 HQCs. PAYG per-HQC price is by quote.	June 2023
 rigetti	Azure Quantum	Rigetti Cepheus-1-108Q	108	Superconducting, tunable qubits, gate-based	7,200 ^b	\$0.02 per 10-ms increment of QPU job execution time. The displayed hourly value is the exact continuous-execution equivalent, not a dedicated reservation.	2026
 Pasqal	OVHcloud Quantum Platform	Pasqal Orion Beta	100	Neutral atoms, analog	3,540	Pay as you go, \$0.983333/s, excluding VAT. Notebook usage is billed separately.	November 2025
 QUANDELA	OVHcloud Quantum Platform	Quandela Belenos-12	12	Photonic, dual-rail / linear optics, gate-based	1,180	Pay as you go, \$0.327777/s, excluding VAT. Notebook usage is billed separately.	April 2026
 IQM	Scaleway QaaS	IQM Emerald, Crystal-54	54	Superconducting transmon, gate-based	3,481 ^a	Dedicated booking at EUR 3,000/h. Shared access is also available at EUR 0.25/circuit + EUR 0.0014/shot.	December 2025
 QUANDELA	Scaleway QaaS	Quandela Belenos, Mosaiq-12	12	Photonic, 12 photons / 24 modes, dual-rail	1,160 ^a	Dedicated booking at EUR 1,000/h. Shared access is also available at EUR 0.30/circuit + EUR 0.000001/shot.	2024
 QAT	Scaleway QaaS	AQT IBEX-Q1	12	Trapped-ion, all-to-all, gate-based	n/a	Shared QPU access at EUR 0.30/circuit + EUR 0.021/shot. No public hourly booking price is currently listed.	2026
 Pasqal	Scaleway QaaS	Pasqal FRESNEL, Orion Beta	100	Neutral atoms, analog	~3,470 ^{a,b}	EUR 3.3/shot at about 0.25 shot/s, listed by Scaleway as about EUR 0.83/s. The hourly number is therefore only an execution-rate equivalent, not a booking tariff.	December 2025
 Infleqtion	Infleqtion QaaS / Superstaq	Sqale	100+ physical	Neutral atoms, gate-based	Not public	Cloud access is metered in shots. Commercialization is still early-stage. The 2025 cloud offering was private beta. Infleqtion reported 12 logical qubits demonstrated by June 2026.	2025, private beta

Table 15: Cloud access prices for selected quantum processors, as of September 12, 2026. Prices are per QPU-hour unless noted and exclude taxes and ancillary classical-cloud costs. ^aConverted at EUR 1 = \$1.16036 on September 12, 2026. Values are rounded to the nearest dollar. ^bContinuous QPU execution-time equivalent. It is not the price of reserving the full machine for one wall-clock hour. Sources: Google [2449, 2450], IBM Quantum [2451], D-Wave Leap [2452, 2453], Amazon Braket [2454–2456] Azure Quantum [2457–2459], OVHcloud [2460–2462], Scaleway [2463–2465], Infleqtion [2466, 2467], OQC [2468, 2469]. © Olivier Ezratty, September 2026.

- **IQM** sells its 5-qubit superconducting QC for less than €1M.
- **Rigetti** sells its 9-qubit Novera superconducting processor for \$900K, but this covers only the hardware below the mixing chamber plate. A complete Novera system including the dilution refrigerator and the control electronics comes closer to \$2.85M, based on the two purchase orders totaling \$5.7M announced in 2025. It also landed an \$8.4M order in January 2026 from India's C-DAC for a 108-qubit chiplet-based system to be installed in Bengaluru in the second half of 2026.
- **PsiQuantum** is building two QCs of 100 logical qubits in Brisbane, Australia and Chicago. The \$1B figure associated with each site comes from funding and public investment commitments. Brisbane alone is backed by A\$940M, about \$620M, of Australian federal and Queensland funding, and both sites are funded by the \$1B series E raised in September 2025. Their announced targets are an operational Brisbane site by the end of 2027 and Chicago in 2028, but with undisclosed capabilities in number of logical qubits, circuit size and clock speed. PsiQuantum will sell computing time on demand when it is made available.
- Various QCs have also been ordered by various public hybrid classical/quantum data centers in Europe and other places.

Based on my knowledge of many vendor roadmaps and pricing plans, it looks like the \$10M level is going to be a key threshold, for both NISQ and FTQC systems, at least in modalities like cold atoms, ions and silicon spin qubits. As we've seen with PsiQuantum, which is replicated with Xanadu, these photonic FTQC QCs may be way more expensive. IBM's future Starling and Blue Jay systems will probably cost several \$10Ms given their size.

5.7.2 Cost structure

Viewed from the vendors, quantum computers have a cost and a price. The ratio between the price and cost is called a "mark up". It builds the operating margin of the vendor, amortizing its R&D, funding its technical support, sales, marketing and other operating costs. And generating some profit, someday. Vendors' markups are not documented publicly and no general factor can be derived from the available disclosures. It may also depend on market concentration, market volume and economies of scale. Markup is also not the same thing as gross margin, and R&D, support and selling costs are not necessarily part of the cost of revenue. For a listed company, a gross margin can be computed from its disclosed revenue and cost of revenue, with the accounting definitions attached.

A quantum computer cost depends on several parameters including its underlying R&D, bill of materials

of off-the-shelf and custom-designed components, manufacturing and integration costs, economies of scale, marketing and sales costs, the cost of maintenance and consumables if any, and finally, the manufacturer's profit. The higher the sales volume, the greater the economies of scale. Volumes are currently very low given most quantum computers are just prototypes that are not yet useful for production-grade applications.

At some point, when and if we reach some quantum advantage threshold, useful applications will emerge. It will first target niche b2b and government markets. Some economists think that quantum computers may offer an economic advantage compared to classical computing even without reaching a computing advantage, thanks to asymmetries in cost structures [2472]. This is still an unproven conjecture since these economists didn't really analyze the real possibility of pre-quantum-advantage NISQ computers to bring any usefulness. The proposed model was only based on economies of scale and the effects of competition.

Then, when applications and innovation ramp-up, we may have a larger number of corporate users. It will justify scaling manufacturing capacities. R&D fixed costs will then be easier to amortize with volume. The cost of goods may also decrease, particularly if technological progress can help get rid of the complicated wirings and electronics that we have today in some of these devices. In a word, the cost also depends on the value, in some virtuous closed-loop circle.

Cost could also be better shared thanks to the emerging practice of open hardware in quantum technologies [2473–2475].

Let's look one by one at the major hardware components of a quantum computer looking at how it will benefit from economies of scale:

- **Control computer(s)**: these are standard rack-mounted servers as well as the associated networking connection. These are the most generic parts of a quantum computer.
- **Chip**: quantum registers chips are the cornerstone of electron-based quantum computers, such as with superconducting and electron spin qubits. Even if they are manufactured in CMOS or similar technologies, their manufacturing volume is very low. Economies of scale are therefore almost non-existent. You don't need such components with cold atoms and trapped ions qubits. It is replaced by specialized optical components to direct the laser beams controlling the qubit atoms. With NV centers, chips can be cheap to manufacture if done in volume.
- **Electronic components**: these are used to create, process, transmit and send quantum gate signals to the qubits. Their technology depends on the type of qubit. These signals are microwaves for superconducting and electron spins qubits, laser-based photons for cold atoms and trapped ions, and some

other varieties of electro-magnetic signals otherwise. Standard and expensive laboratory equipment is used for microwave generations such as those from Rohde & Schwarz. More integrated equipment is sold by companies like Zurich Instruments, Qblox, Quantum Machines and Keysight. It is using customized FPGA and rather standard electronic components. When these tools are miniaturized as ASICs running at room temperature or cryo-CMOS running at temperatures below 4K, their small economies of scale make it rather expensive.

- **Cabling:** niobium-titanium superconducting cabling used to feed superconducting and electron spin qubits with microwaves is expensive, on the order of \$3K per line, a figure that depends on the supplier, the specification and the volume. The number of lines per qubit depends on the configuration. A flux-tunable transmon architecture uses drive, flux and readout lines, and readout multiplexing brings the practical average closer to two lines per qubit in current designs. This drives high costs for manufacturing superconducting and electron-spin based qubits systems. The main companies providing these cables are Coax&Co and Delft Circuits.
- **Cryogenics:** these are standard systems but marketed in low volumes. They can cost up to \$1M for superconducting and silicon qubits. Their cost is about one order of magnitude lower for the cryostats cooling photonic components and their superconducting nanowire detectors, which operate at 4K or below 1K and not at 15 mK. Large cryostats use an enclosed cooled system with many cylindrical layers of protection, a GHS (gas handling system), a compressor (such as those coming from CryoMech and Sumitomo) and another compressor used to cool the water feeding the primary compressor.
- **Lasers** used in trapped ion, cold atom and photonic QCs can be costly depending on their characteristics and power. There are many such lasers in cold atom QCs. In all these systems, these are fragile components that break from time to time. QC vendors must have spare units in-store to replace them quickly when they fail.
- **Consumables:** in quantum computers operating at very low temperatures, there is at least some liquid nitrogen, used in the cold traps of the gas handling system and for pre-cooling, plus gaseous helium 3 and 4. Dry dilution systems are cryogen-free in steady state and their helium mixture circulates in a closed loop, so it is not consumed, but the initial helium 3 charge and the top-ups needed after leaks remain expensive.
- **Casing:** this is about steel, glass and design with some specifics linked to vibrations isolation.
- **HPC or supercomputer:** many quantum algorithms, particularly in quantum chemistry, will re-

quire some heavy classical computing for pre and post-processing. Since there are only a handful of exascale supercomputers in the world, a few hundred TOP500-class systems, and they are quite costly, it may become an economical barrier to using quantum computers at large scale. It will be interesting to see if quantum computing use case related to various industry day-to-day operations can accommodate lighter classical HPC to drive their QCs. This is one of the reasons why Nvidia is highly engaged in this domain. They will sell DGX servers (using 8 GPGPU each) like hotcakes to quantum computer industry vendors. These systems will handle many tasks, the most important being the classical part of error correction (error syndrome decoding).

- **SLA and MTBF.** These are two related concepts that are not commonplace in the current quantum computing landscape. “Service level agreements” are vendors engagement on the quality of service they deliver around their QCs. For example, what is their percentage of availability time. What are the contractual engagements and penalties to restore a malfunctioning QC. This is connected to the “mean time between failure” of the QC and its components. For example, we know that the lasers used in atom and photon systems can easily break. How fast are they replaced? This will drive some incurred costs for the vendors which will have to keep an inventory of costly spare parts. QuEra mentions this notion of SLA in a 2025 document on QC commercial reality milestones [2476]. D-Wave has SLAs in place since 2024 and advertises 99.9% availability and uptime for its Leap cloud service.

As quantum technologies mature, some cost structures will increase, and others will decrease. Economies of scale will do the rest. Comparisons will also be key. We can also expect that quantum computers will have a price constrained by the most expensive supercomputer prices given the available budgets of their users, mostly public research and large corporation organizations. In practice, many quantum computers will be usable as resources in the cloud and at a relatively moderate cost, although currently much higher than classical HPC cloud resources. This is what IBM, Rigetti, D-Wave, Microsoft and AWS (with third-party machines for the latter two) are already offering. Microsoft and AWS quantum cloud pricing is already quite high. Then, one can wonder about its added value.

Some early estimates of the cost of FTQC systems supporting thousands of logical qubits have been made by companies like Google. Scaling today’s laboratory costs linearly with the required physical qubits lands in the \$B range, but that extrapolation is not a valid cost model for a future system. Component costs, multiplexing, integration and supplier volumes will all change with scale.

In a white paper published in April 2025, Matthew Gould and Kate Prebble from the New Zealand investment fund ICM laid out interesting views on the economics of quantum computers [2477]. They defined the economic conditions of FTQC viability and conveyed the message that the difference in capex and opex between the qubit types competing with each other will become a key differentiating factor. This is in line with what's in this book, particularly in the energetics section and in the chart in Figure 298 on power estimates (used in their paper using an earlier version, although their own Table 6 energy estimates are way too pessimistic in the case of IBM/Google, trapped ions and cold atoms). Their rationale was also made to justify their investment in Diraq (Australia), with a qubit type that, if it works, can be highly scalable and cost efficient. This is a good thing to start this debate!

How about quantum computers becoming consumer products and having their price decrease to a couple thousand dollars? One should never say never, of course, but this is quite unlikely based on our current knowledge of the technology [2478]. Whatever the qubit modality, it requires components that are not as easy to miniaturize as the chips that are in your smartphone and laptop. You also need cryogenics. At last, it doesn't necessarily make sense from the applications viewpoint. What quantum applications would you really need on a personal device that can't be implemented with classical computing and/or some cloud access? It is your call to invent one!

5.7.3 Unit sales and deployments

Most quantum computing startups are obliged to provide potential investors with sales forecasts. These often include unit sales dispatched to end-user customers and to cloud providers. These are gut feeling-based and not necessarily based on a solid rationale. It is not just a matter of market size and competition. It is related to two unknowns at this point: the real value created by these future QCs and their cost structure. The value depends on algorithmic progress that is not there yet, particularly to unlock the larger business market with optimization applications, including those which may be quantum machine learning based.

We can however use comparables. As seen in a previous subsection on the energetics of quantum computers, we can have a look at the size of the supercomputing market and use their price and energetics as indicators. What it shows is that the price and energetics are good indicators of a potential volume market. If QCs are sold at between \$50M and \$500M, their market will be limited to a hundred systems. If however, FTQC QCs are available at below \$10M, then the market might be larger, up to a thousand units. And if it goes as low as \$1M, then, we could envision a much larger market, pending solving the business algorithm problem.

Another question relates to the share of FTQC QCs that will be sold to end-users ("on premise") and to cloud vendors. Defense and government related end-users will probably acquire these QCs while business users may select to do it if they have applications requiring exclusive availability for these QCs. Again, this will depend on the viable applications running on FTQC QCs and bringing some solid quantum advantage compared to best-in-class classical solutions.

We can expect these FTQC QCs to be replaced every 5 years, as are most supercomputers. This will come from a natural obsolescence of QCs and to the way it will be amortized by either end-users or cloud vendors.

5.7.4 Benefits

We must make an economic distinction between **cost** (of R&D, goods and manufacturing), **price** (how much is it sold or rented) and **value** (what value is it bringing to customers, particularly, compared with existing classical computing solutions).

Right now, the equation is simple: costs are quite high, prices are high as well when computers are sold on premise (particularly those using superconducting qubits) and value is low at this point and is positioned in the educational and proof-of-concept realms. Some interesting case studies at the edge of quantum advantage have been documented, but most of the time, to solve condensed matter physics problems with no known commercial value. The value there is for the benefit of science, and is indirect.

Some economists are trying to quantify the economic advantage of quantum computers even without them bringing some quantum computing advantage. Their assumptions are often quite sketchy without a good understanding of the underlying use cases, as already seen above with the quantum economic advantage model [2472].

Other economists from the University of Cambridge and Bandung Institute of Technology in Indonesia tried in 2023 to evaluate the productivity gains or losses generated by quantum computers adoption [2479]. Their conclusion was rather negative, like what supposedly happened with the advent of classical computing in the 1970s and 1980s with a productivity paradox coming from the initial need to invest in new equipment and know-how. These costs could be higher with quantum computing for various reasons, one being the difficulty of finding relevant use cases and the other being the threats generated by quantum computers on cybersecurity.

The experts then proposed to demonstrate the value of quantum computers for societal challenges. Why not, but these are not necessarily aligned with the corporation goals whatever they say in their CSR (Corporate Sustainability Responsibility) reports. Then, they advise the creation of a unified and standardized language to identify the right problems for quantum computers

and prepare data in a quantum-ready format. Finally, they recommend deploying either PQC or QKD cryptography. But they have no idea of the cost structure of quantum computing and how the industry would consolidate around standards and how its size would generate economies of scale.

A joint team from the MIT FutureTech Lab and Accenture developed its own quantum economic advantage calculator, to estimate how and when quantum computing would bring an economic advantage [2480]. It applied this method to estimate when quantum chemistry would deliver some economic advantage, with focusing on ground state estimations [2481]. Under their assumptions, the model outputs 2031 for an FCI-equivalent QPE chemical simulation and a date beyond 2050 for matching DFT, the cheapest classical method and therefore the hardest one to beat (Table 16). Their model assumes an overhead of about 10^{13} in favor of classical hardware at comparable cost, which is what the algorithmic advantage has to overcome. But all these estimations are based on various assumptions that deserve scrutiny like the real cost of classical simulations and the overhead of quantum error correction.

5.7.5 Ecosystem structure

The quantum computer market is currently highly fragmented due to the sheer variety of qubit types and subtypes. We have over 90 vendors and at least about 20 with “commercial” systems in the NISQ realm. About 225 QC units have been sold or deployed so far according to my own data consolidation as of September 2026 [2482], including over 40 retired IBM systems, which leaves about 185 in operation (Figure 313 and Figure 312). That total aggregates systems of differ-

ent nature: units sold to a customer, systems installed on customer premises, vendor-owned systems exposed through a cloud service, upgrades of an existing system, retired systems and contracts not yet delivered. Geographic allocation is ambiguous for vendor-owned cloud systems, which are counted in the vendor’s country rather than in the users’ countries.

There are of course lots of unknowns on which company and technology will scale, when and at what cost. What will defragment this market? Usually, it comes from a mix of technology, platform, and ecosystem leadership and economies of scale. It explains the past and present dominance of the IBM 360, MS-DOS, Windows, Intel CPUs in PCs and server, Android and iOS on smartphones, Google, Facebook and Amazon, TensorFlow and PyTorch, ChatGPT and other Nvidia of this world.

Here, the technology is not settled, nor its cost-benefit structure. Still, some key segments of the ecosystem are being defragmented: in software development (Qiskit), control electronics (Quantum Machines), cryogenics (Bluefors), lasers (Toptica), and software engineering (Classiq, Riverlane). It looks a bit like the PC industry, with stacks of specialized players capturing a dominant position (>50% market share) in their domain.

Major IT companies such as IBM, Google, Intel, Amazon and Microsoft can fund quantum computing R&D investments with a very long-term vision. They have profitability, cash and the ability to attract skills to do so. You may still notice that these large IT vendors have not yet engaged in a startup acquisition spree like they did in the fields of artificial intelligence and other emerging technologies. They are still hedging their bets with collaborating with multiple partners, like Microsoft with Quantinuum and Atom Computing and investing in PsiQuantum and Photonic Inc.




Figure 312: Consolidation of quantum computer unit sales or deployments in customer premises or in cloud offerings as of September 2026 per provider and acquirer country. The USA total contains all past and existing IBM QC put online in their cloud service, given the majority of these systems have been retired (>40 over 6 in production). The chart showcases the position of the European Union member states, next to the UK. © Olivier Ezratty data, September 2026. Sankey chart realized with Python code generated by ChatGPT based on the author's QPU sales and deployment tracking data.





Figure 313: Consolidation of quantum computer unit sales or deployments as of September 2026 per provider region. The chart accounts for past and present QCs used by IBM and D-Wave in their cloud services. It doesn't account for the R&D contracts from DARPA QBI in the USA and ProQCima in France, which cannot yet be considered as QC orders, and for non-commercial QCs (like the QuTech QCs used in the Quantum Inspire online service in the Netherlands). The consolidated data is based on public announcements and may not be complete, particularly with regards to quantum systems deployed in China. © Olivier Ezratty data, September 2026.

Classical Method	Classical Time	$O(N^3/\epsilon)$ QPE	$O(N^2/\epsilon)$ QPE
Density Functional Theory (DFT)	$O(N^3)$	N/A	> 2050
Hartree Fock (HF)	$O(N^4)$	> 2050	2044
Møller-Plesset Second Order (MP2)	$O(N^5)$	> 2050	2038
Coupled Cluster Singles and Doubles (CCSD)	$O(N^6)$	2044	2036
Coupled Cluster Singles and Doubles Perturbative Triples (CCSD(T))	$O(N^7)$	2036	2034
Full Configuration Interaction (FCI)	$O^*(4^N)$	2032	2031

Table 16: Comparison of classical and quantum (QPE) methods with their asymptotic times and the first year of quantum economic advantage produced by the cited model. The years are outputs of that scenario. The scalings shown are asymptotic in N and should be read as exponents, $O(N^4)$ and $O^*(4^N)$ for instance. The physical qubit error rate is the usual $\epsilon = 10^{-3}$ that shows up in many resource estimates, N representing the number of relevant basis functions. The last column used more efficient circuit in computing time, at the expense of a larger number of logical qubits. Source: [2481].



5.7.6 Bibliography and notes

These are the bibliographical references and notes for the **Economics** subsection.

- [2449] [Access and authentication](#), by Google Quantum AI, July 2026.
- [2450] [Willow Early Access Program](#), by Google Quantum AI, March 2026.
- [2451] [IBM Quantum Computing: Products and services](#), by IBM Quantum, 2026.
- [2452] [Leap Service](#), by D-Wave Quantum, 2026.
- [2453] [Operation and Timing](#), by D-Wave Quantum, 2026.
- [2454] [Amazon Braket Pricing](#), by Amazon Web Services, 2026.
- [2455] [Amazon Braket supported regions and devices](#), by Amazon Web Services, 2026.
- [2456] [Working with reservations](#), by Amazon Web Services, 2026.
- [2457] [Pricing Plans for Azure Quantum Providers](#), by Microsoft, December 2025.
- [2458] [Quantinuum provider](#), by Microsoft, May 2026.
- [2459] [Rigetti provider](#), by Microsoft, 2026.
- [2460] [Price list: A comparison of our Public Cloud offers](#), by OVHcloud, 2026.
- [2461] [OVHcloud unveils its Quantum Platform: the first European solution for accessing quantum computers in the cloud](#), by OVHcloud, November 2025.
- [2462] [OVHcloud accelerates the European quantum industry with Quandela's Belenos](#), by OVHcloud, April 2026.
- [2463] [Quantum as a Service](#), by Scaleway, 2026.
- [2464] [IQM QPUs Information](#), by Scaleway, 2026.
- [2465] [AQT and Scaleway partner to deliver Quantum Computing to Europe](#), by Scaleway, February 2026.
- [2466] [Sqale Quantum Computer](#), by Infleqtion, 2026.
- [2467] [Annual Report for the fiscal year ended December 31, 2025](#), by Infleqtion, March 2026.
- [2468] [OQC Cloud](#), by OQC, 2026.
- [2469] [OQC launches OQC Toshiko, the world's first enterprise ready quantum platform](#), by OQC, November 2023.
- [2470] [Advancing European Quantum Computing: Signature of the Procurement Contract for EuroQCS-Poland](#), The European High Performance Computing Joint Undertaking (EuroHPC JU), July 2024.
- [2471] [Signature of the Procurement Contract for the EuroHPC Quantum Computer Located in Italy](#), by EuroHPC JU, The European High Performance Computing Joint Undertaking (EuroHPC JU), March 2025.
- [2472] [Quantum Economic Advantage](#), by Francesco Bova, Avi Goldfarb, and Roger G. Melko, Management Science, February 2023 (11 pages).
- [2473] [Open hardware solutions in quantum technology](#), by Nathan Shammah, Irfan Siddiqi, William J. Zeng, et al., APL Quantum, March 2024 (17 pages).
- [2474] [The Economics of an Open-Source Quantum Computer](#), by Francesco Bova and Roger G. Melko, arXiv, January 2025 (11 pages).
- [2475] [An open-source initiative would benefit quantum computing](#), by Francesco Bova and Roger G. Melko, Nature Reviews Physics, August 2025 (2 pages)  (preprint on [arXiv](#)).
- [2476] [Three Steps \(and nine milestones\) to Commercial Reality in Quantum Computing](#), by QuEra, September 2025 (4 pages).
- [2477] [An economic analysis of quantum computing](#), by Matthew Gould and Kate Prebble, ICM, April 2025 (30 pages).
- [2478] [Will Quantum Computing Ever Become a Consumer Technology?](#), by Victor Kerros, Quantum Times, July 2026.
- [2479] [How to introduce quantum computers without slowing economic growth](#), by Chander Velu and Fathiro H. R. Putra, Nature, July 2023.
- [2480] [Introducing the Quantum Economic Advantage Online Calculator](#), by Frederick Mejia, Hans Gundlach, Jayson Lynch, Carl Dukatz, Andrew Lucas, Eleanor Crane, Prashant Shukla, and Neil Thompson, arXiv, August 2025 (11 pages).
- [2481] [Quantum Advantage in Computational Chemistry?](#), by Hans Gundlach, Keeper Sharkey, Jayson Lynch, Victoria Hazoglou, Kung-Chuan Hsu, Carl Dukatz, Eleanor Crane, Karin Walczyk, Marcin Bodziak, Johannes Galatsanos-Dueck, and Neil Thompson, 2025 IEEE International Conference on Quantum Computing and Engineering (QCE), September 2025 (11 pages)  (preprint on [arXiv](#)).
- [2482] [QPUs Worldwide](#), by Neil Abroug, Github, April 2025 (>100 entries).

This **Economics** subsection contains 34 bibliographical references and notes containing at least 97 pages.

5.8 Quantum computing engineering key takeaways

- A quantum computer uses physical qubits of different nature, the main ones being superconducting qubits, electron spin qubits, cold atoms, trapped ions and photons. They all have pros and cons, and no one is perfect at this stage. Future systems may combine several of these technologies.
- Many key parameters are required to create a functional quantum computer. It must rely on two-level quantum systems used as qubits. These must be initializable and manipulable with a universal gate set that can approximate any unitary transformation on the relevant qubits to the desired accuracy. Qubits must be measurable when required by the algorithm. More general quantum channels are not arbitrary gate operations and require open-system dynamics or constructions involving ancillas, measurements or discarding subsystems. Coherence times and physical error rates must be compatible with the required circuit depth and error-management scheme.
- Most quantum computers are composed of several parts: a qubit chip (for solid-state qubits but also for trapped ions), vacuum enclosures (particularly with cold atoms and trapped ions) or sources, waveguides and detectors (photons), usually housed in a cryogenic vacuum chamber, some electromagnetic waves conveyed to the qubits with laser beams, or coaxial cables guided microwaves or direct currents, and a classical computer driving these electronic and photonic components. Qubits are driven by electromagnetic waves in various parts of the spectrum: UV, optical photons, infrared photons, microwave and radio-frequency signals.
- Errors are unwanted effects happening with qubits during initialization, gate operations, idle times, and readouts which affect the quality of the results. Noise is the main type of error source of stochastic nature that can be endogenous or exogenous. Decoherence progressively destroys superposition and entanglement in, and between qubits. Scientists devised quantum error mitigation (after computing) and active correction (during computing) schemes to address these problems. Quantum error correction relies on creating logical qubits composed of many physical qubits, the number depending on the quantum circuit size and many other parameters. A logical qubit is not a “perfect” qubit. It has a residual error rate designed to support the quantum circuit size. The number of physical qubits per logical qubit depends on it, the error correction codes, the qubits’ physical connectivity and the physical qubit error rates. This creates huge scalability challenges, many of them with classical qubit control technologies. Fault-tolerant quantum computing (FTQC) is not just about using logical qubits. Error correction must happen faster than error creation and errors must not propagate within the quantum computer when logical gates are implemented.
- Many quantum circuits also require data preparation or storage, for example in quantum machine learning or oracle-based algorithms such as Grover search. Quantum-state memories exist in experimental and limited practical forms. They should be distinguished from qRAM, which aims at coherent random access to a large classical or quantum-addressed data set, and from circuit-based QROM. Large-scale fault-tolerant qRAM with efficient coherent random access is not currently available.
- The energetic cost of quantum computing is both a potential benefit but also an immense challenge, particularly when a large number of physical qubits are required to create large scale fault-tolerant computers. All components must be carefully designed to take into account the cryogenic cooling power, control electronics, lasers, cabling as well as the available space to house cabling and cryo-electronics. Based on vendor roadmaps, there seems to be a 1 to 100 scale of power consumption for future FTQC quantum computers. This explains the creation of the Quantum Energy Initiative in 2022, which gathers a community of researchers and industry vendors and organizations working collectively on this topic.
- The book describes the Lieb-Robinson bounds, which quantifies how fast information and entanglement propagate in a quantum system, and its potential impact on quantum-computing systems architectures from hardware to compilers and error correction.
- The economics of quantum computers are still uncertain due to their immaturity and the current low manufacturing volume. Uncertainty is also strong with regard to the feasibility of scalable quantum computers. The scalability challenges ahead are enormous. One of them is to benefit from actual algorithm speedups when including all end-to-end computing operations. Another one is about expanding the limits in our ability to control a large number of entangled quantum objects, and with excellent fidelity.

Page intentionally left blank.

Page intentionally left blank.

Page intentionally left blank.

