Introduction

amU QuantA
Aix Marseille Université

the challenges of FTQC

olivier ezratty
olivier@oezratty.net www.oezratty.net

Marseille, November 26th 2026

1

This talk is about the challenges of fault-tolerant quantum computing (FTQC). I would like to explore how hardware and software come together, from the mathematical problems we want to solve to the physical systems required to solve them. Quantum computing is a particularly interesting field for a generalist engineer because it brings together many scientific and technological disciplines.

agenda

- **from use cases to algorithms**
- **FTQC resource estimates**
- **FTQC hardware roadmaps and challenges**
- **quantum computing energetics**
- **quantum computing ecosystems**

I will start with applications and algorithms, then move to resource estimates, quantum error correction (QEC), fault tolerance and hardware roadmaps. I will devote some time to quantum computing energetics, one of my main areas of activity as a co-founder of the Quantum Energy Initiative (QEI) along with Alexia Auffèves and Robert Whitney, two CNRS researchers located respectively in Singapore and Grenoble. Finally, I will discuss the ecosystem and the way generative AI is changing scientific research.

From use cases to quantum algorithms

from use cases to quantum algorithms

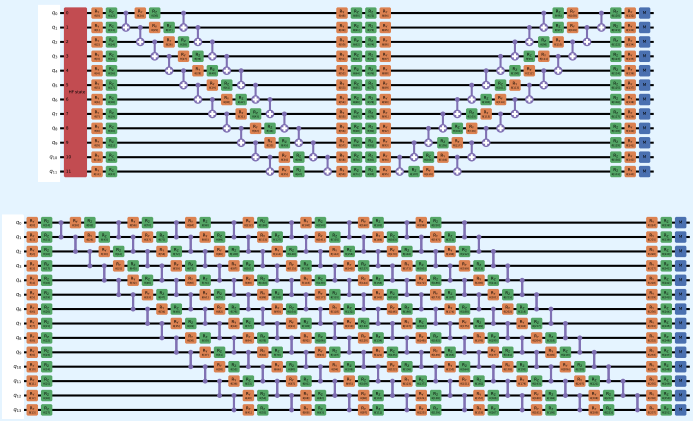

typical difficult problems

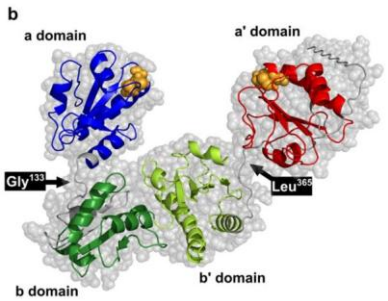

$$i\hbar\frac{\partial\Psi(x,t)}{\partial t} = -\frac{\hbar^2}{2m}\frac{\partial^2\Psi(x,t)}{\partial x^2} + V(x)\Psi(x,t)$$

solving Schrodinger's wave equation to simulate quantum matter

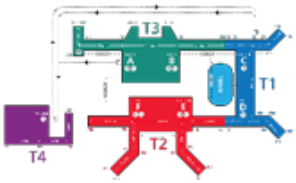

combinatorial optimizations



solving partial differential equations

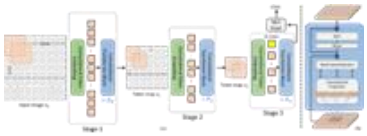

machine learning and deep learning

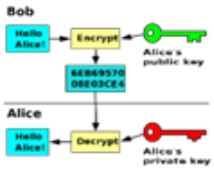

breaking asymmetric cryptography keys

4

Why do we need quantum computers? Richard Feynman's proposal in 1981 to use quantum systems to simulate other quantum systems is still highly relevant. When we look at the practical applications of gate-based quantum computing and analog quantum simulation, the simulation of quantum matter remains one of the most convincing application areas.

Other frequently discussed applications include combinatorial optimization, machine learning and the resolution of partial differential equations (PDEs). The traveling salesperson problem (TSP) is often used as an illustration of optimization, but it is a poor example of practical quantum advantage because classical heuristics are already very efficient. More constrained problems, such as the capacitated vehicle routing problem (CVRP), are harder, although this does not establish a quantum advantage.

Quantum machine learning is another active research area, but we have yet to demonstrate that a quantum computer can efficiently execute large-scale machine learning or deep learning workloads. Solving PDEs, including problems related to the Navier–Stokes equations, is somewhere in between. Some quantum algorithmic primitives are promising, but data loading, nonlinearities and output extraction are major obstacles.

Cryptanalysis is a special case. Shor's integer factoring algorithm provides a very clear theoretical motivation for building a fault-tolerant computer. Its exponential improvement over the best known classical factoring methods is often cited, although no exponential lower bound for classical factoring has been proved. Breaking RSA-2048 is not an application that creates conventional industrial value, but it matters to governments and to the security of the internet. It is somewhat paradoxical that one of our clearest examples of a quantum computational advantage is also a threat to existing cryptography.

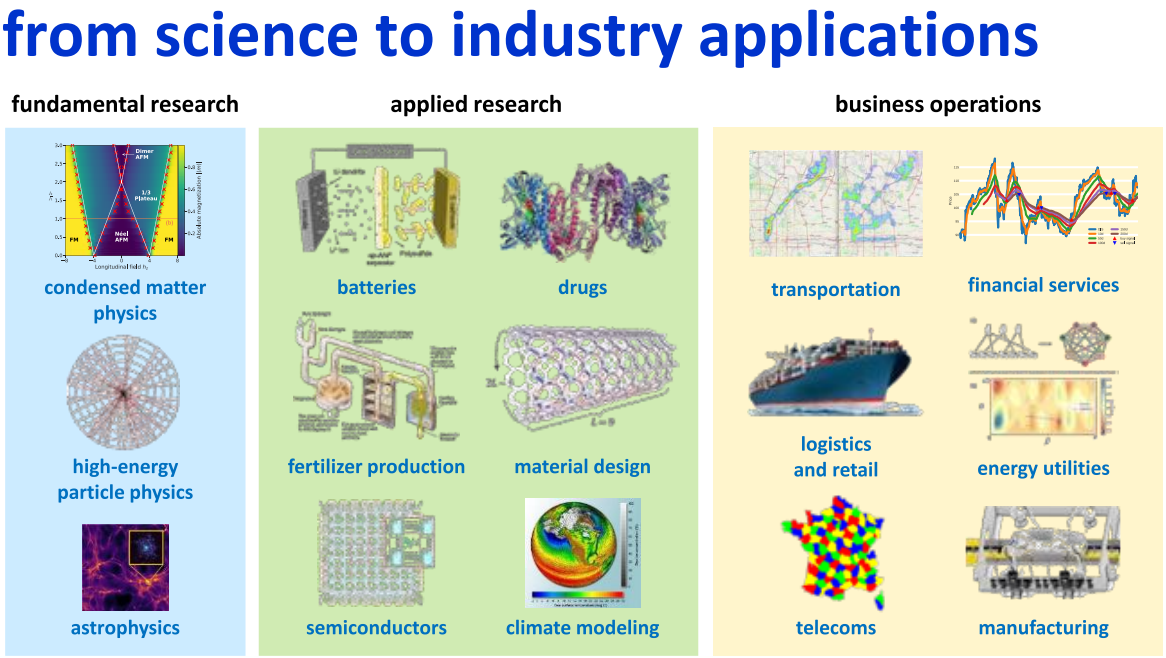

Who would use these applications? I distinguish three categories. Fundamental research includes condensed matter physics, high-energy physics and astrophysics. Researchers working at CERN or analyzing astronomical observations do not necessarily have an immediate commercial application in mind. I had the opportunity to discuss quantum computing applications at CERN in January 2025.

Applied research lies at the intersection of science and industry. It includes the development of new fertilizers, drugs, batteries, materials and semiconductor technologies. Quantum computers could contribute to creating digital twins at the atomic and electronic scales. Once a new material or chemical process has been identified, its industrialization remains largely classical.

The third category comprises the day-to-day operations of businesses: transportation, logistics, financial services, manufacturing, telecommunications and energy utilities. Matching electricity demand and supply in a grid with intermittent renewable production is an example of a potentially valuable optimization problem. These industries already have classical solutions, but they would like to improve their efficiency.

There is a paradox at the intersection of science and economics. As we move from fundamental research toward business operations, the potential market increases, but identifying algorithms with a useful quantum advantage becomes harder. The largest markets require optimization and machine learning, precisely the domains where the case for quantum advantage remains particularly uncertain. This differs from generative AI, which can be used across most professions without such a pronounced technological

distinction between market segments. Investors often dream about the largest market, which helps explain some of the hype around quantum computing.

typical target Hamiltonians per paradigm

gates-based quantum computing

Heisenberg model (VQE)
spin dynamics, quantum magnetism, superconductivity.

$$H = \sum_{\langle i,j \rangle} \left(J_x S_i^x S_j^x + J_y S_i^y S_j^y + J_z S_i^z S_j^z \right) + \sum_i h_i S_i^z$$

electronic structure model (VQE, QPE)
encoded into qubit gates using Jordan–Wigner and Bravyi–Kitaev transforms, also applicable with HEP.

$$H = \sum_{p,q} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{p,q,r,s} g_{pqrs} a_p^\dagger a_q^\dagger a_r a_s$$

Fermi-Hubbard model (VQE, QPE)
strongly correlated systems in condensed matter physics, magnetism, Mott insulators, high T_c superc.

$$H = -t \sum_{\langle i,j \rangle, \sigma} \left(\hat{c}_{i,\sigma}^\dagger \hat{c}_{j,\sigma} + \text{h.c.} \right) + U \sum_i \hat{n}_{i,\uparrow} \hat{n}_{i,\downarrow}$$

create unitary U based on H so that $||U - e^{-iHt}|| < \epsilon$ (error rate), and find eigenvalue or eigenstate of H of interest

analog quantum simulation

XX Ising models
quantum transport, superfluid-Mott insulator transitions, topological phases...

$$H = \sum_{i<j}^{N} J_{ij} \, \hat{\sigma}_i^x \hat{\sigma}_j^x$$

XY Ising models
non-trivial topological effects, real-time quench dynamics...

$$H = \sum_{i<j}^{N} J_{ij} \, (\hat{\sigma}_i^x \hat{\sigma}_j^x + \hat{\sigma}_i^y \hat{\sigma}_j^y)$$

XXZ Ising models
ferromagnetic, antiferromagnetic, spin-liquid states, U(1) LGT simulations...

$$H = \frac{1}{2} \sum_{i \neq j} J_{ij}^x (\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y) + J_{ij}^z \sigma_i^z \sigma_j^z$$

find σ_i^x, σ_i^y, and/or σ_i^z minimizing a cost function that is the minimum eigenvalue of H

quantum annealing

Ising Hamiltonian
spin–spin interactions, quantum magnetism, $\mathbb{Z}_2$ and $U(1)$ lattice models.

$$H = \sum_i h_i \sigma_i^z + \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z$$

QUBO formulation
optimization problems, graph partitioning, route planning

$$H = \sum_i a_i x_i + \sum_{i,j} b_{ij} x_i x_j$$

MaxCut problems
solving various graph problems.

$$H = \sum_{\langle i,j \rangle} w_{ij} \left(\frac{1 - \sigma_i^z \sigma_j^z}{2} \right)$$

find σ_i^z minimizing a cost function that is the minimum eigenvalue of H

(cc) Olivier Ezratty, 2025

6

We also need to distinguish the different computing paradigms. Gate-based quantum computing can implement models such as the Heisenberg Hamiltonian, electronic structure Hamiltonians and the Fermi–Hubbard model. Analog quantum simulators can implement various spin models, while quantum annealers typically implement Ising Hamiltonians and quadratic unconstrained binary optimization (QUBO) formulations.

Theoretical polynomial relationships between some computing paradigms do not necessarily imply practical equivalence. A polynomial overhead can be substantial. The hardware constraints and the accessible Hamiltonians matter. In this talk, I mainly focus on gate-based FTQC.

Hamiltonian → unitary → circuit

H

problem Hamiltonian describing a system energy

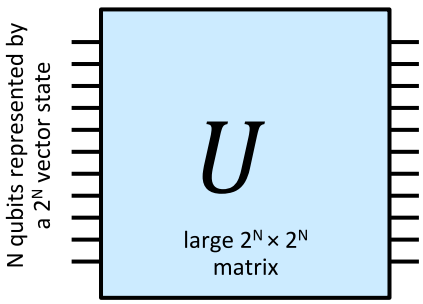

we are looking for U that encode the Hamiltonian variables into qubits and approximates H with:

$$||U - e^{-iHt}|| < \epsilon$$

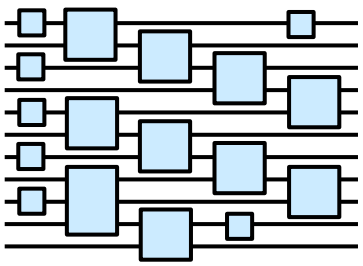

U is decomposed into a series of quantum gates to find eigenvalues or eigenstates of interest of H, i.e. a ground state

7

For many quantum simulation problems, we start from a Hamiltonian describing the system. We then need to construct a unitary that approximates its evolution, or another suitable encoding of the Hamiltonian, and decompose that unitary into elementary quantum gates. This is not automatic. Transforming a problem Hamiltonian into an efficient circuit is itself a computational problem, much of which is solved classically before the quantum computer executes anything.

what is a quantum algorithm?

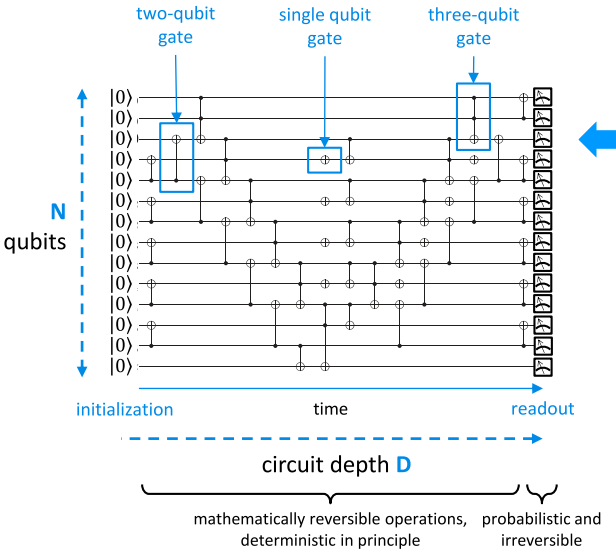

quantum algorithm: solves a mathematical problem with generating and executing a quantum circuit.

quantum circuit: series of operations (gates) acting on a quantum memory, on individual qubits or with connecting them.

matrix computing: linear operations on an internal memory of 2^N complex numbers.

mechanism: uses constructive and destructive interferences in the state vector to yield a value of interest as a series of classical bits.

classical computing: used to prepare the circuit, encode the data, and extract some useful result from multiple circuit runs and measurements.

8

What is a quantum algorithm? A quantum circuit is a sequence of operations on qubits. Single-qubit gates manipulate individual qubits, two-qubit gates can generate entanglement, and three-qubit gates such as the Toffoli gate are useful for reversible arithmetic, notably in Shor's algorithm.

I distinguish the circuit from the algorithm. The latter encompasses the classical techniques required to generate the circuit, prepare the data, execute the circuit and extract the result. We discussed this distinction in the context of a report from the French Academy of Technologies.

Mathematically, the unitary part of the circuit applies matrices to a state vector. The mechanism behind quantum computing is often described as parallelism. I used to think about it that way when I started studying the field eight years ago. However, the useful mechanism is the construction of interference patterns that amplify or suppress amplitudes associated with particular outcomes. We cannot simply retrieve all the values in the state vector.

The unitary evolution is theoretically deterministic and reversible, but the state vector is not directly observable. Measurement produces classical bits and is generally probabilistic and irreversible. Consequently, the number of times we need to execute a circuit to extract a useful result is an essential part of its cost.

how **Quantum Phase Estimation** works

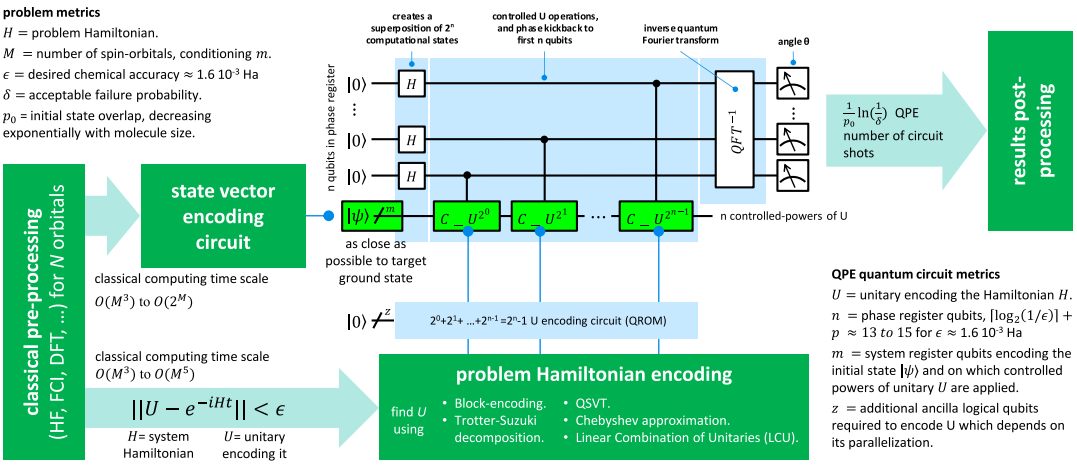

inspired by Quantum chemistry, classical heuristics, and quantum advantage by Garnet Kin-Lic Chan, arXiv, July 2024 & On the feasibility of performing quantum chemistry calculations on quantum computers by Thibaud Louvet, Thomas Ayral, and Xavier Waintal, arXiv, June 2023-October 2024.

9

Quantum Phase Estimation (QPE) illustrates these issues. In chemistry, we may want to estimate the ground-state energy of a molecule, an excited-state energy or another eigenvalue. The circuit uses a system register initialized in a state with sufficient overlap with the target eigenstate, controlled powers of a unitary encoding the problem and an inverse Quantum Fourier Transform (QFT) on a phase register.

There is a double challenge. We need to prepare a suitable initial state, often using classical chemistry methods, and encode the Hamiltonian into the unitary operations. Techniques such as Trotter–Suzuki decomposition, block encoding, linear combinations of unitaries and qubitization can be used depending on the algorithm. The circuit is a combination of primitives, each with its own cost.

The precision of the phase estimation and the overlap with the desired eigenstate affect the resources. If the probability of successfully sampling the desired eigenstate is p_o, repeated independent runs on the order of $\ln(1/\delta)/p_o$ may be needed to obtain a result with failure probability δ, assuming the phase estimation itself succeeds. Chemical accuracy is commonly associated with approximately 1.6×10^{-3} hartree. The classical preprocessing, quantum circuit execution and classical post-processing must all be included in the resource estimate.

5-qubit QFT unitary

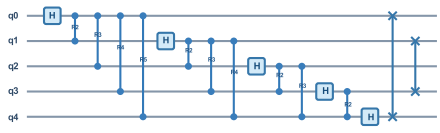

Rk = controlled phase P(2π/2^k). Final SWAPs reverse the qubit order.

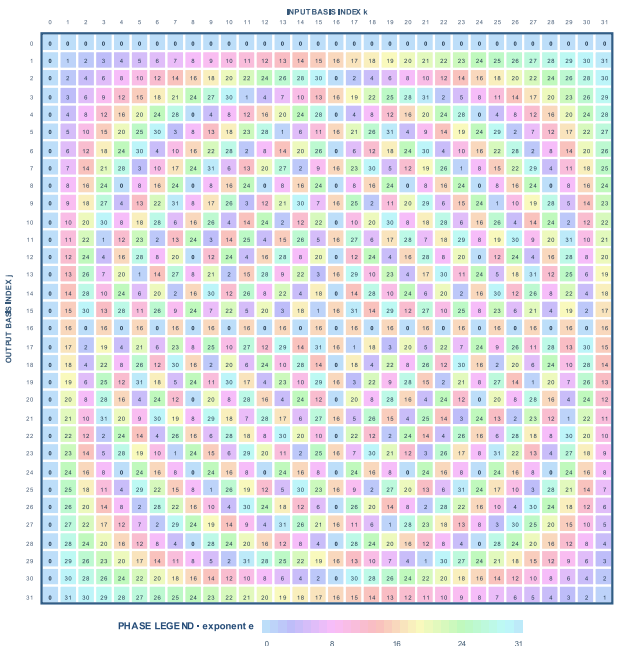

provides some insights on why circuits are useful

- they turn a 17-gate circuit into a unitary matrix with 961 non-zero values.
- modulo a symmetry around the diagonal due to the QFT unitary being equal to its transpose unitary.
- imagine the relative sizing with >50 qubits!

As a software engineer, I like visualizing data. I recently used ChatGPT to generate visualizations of the QFT, starting with five qubits. The circuit has 17 gates when controlled rotations and final SWAPs are counted as elementary gates. Its unitary is a 32 × 32 matrix with 961 nonzero coefficients, all of identical magnitude. The colors encode their phases. A relatively compact circuit can therefore describe a large dense matrix.

7-qubit QFT

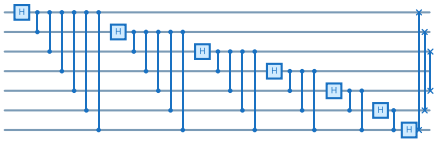

two additional qubits…

- turns a 31-gate circuit into a unitary matrix with 16,129 non-zero values.



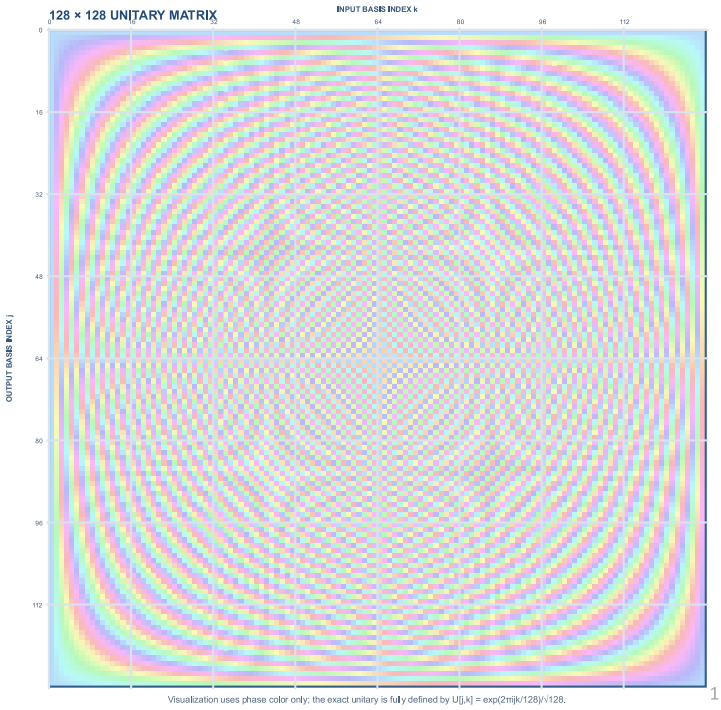

With seven qubits, the circuit has 31 gates and the matrix has 128 × 128 = 16,129 coefficients. The difference in scaling is striking: the standard QFT circuit grows quadratically with the number of qubits, whereas the explicit matrix grows as 4 to the power of that number. This illustrates why decomposing the unitary into gates is useful, although it does not by itself establish an end-to-end computational advantage.

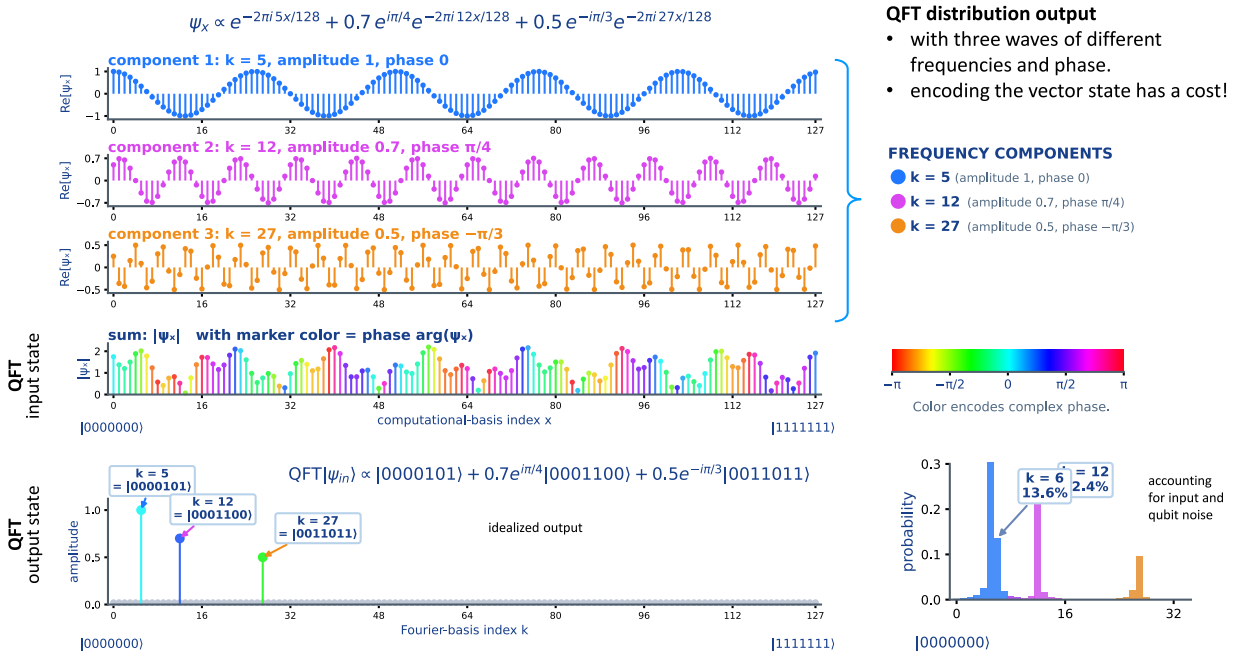

QFT distribution output

- with three waves of different frequencies and phase.
- encoding the vector state has a cost!

FREQUENCY COMPONENTS

- **k = 5** (amplitude 1, phase 0)
- **k = 12** (amplitude 0.7, phase π/4)
- **k = 27** (amplitude 0.5, phase −π/3)

12

For a seven-qubit state encoding three waves with different frequencies, amplitudes and phases, the ideal QFT concentrates the output on the three corresponding frequency components. Noise spreads some probability to other outcomes. Preparing the input state also has a cost, and extracting the phase information requires more than a simple computational-basis measurement.

10-qubit QFT

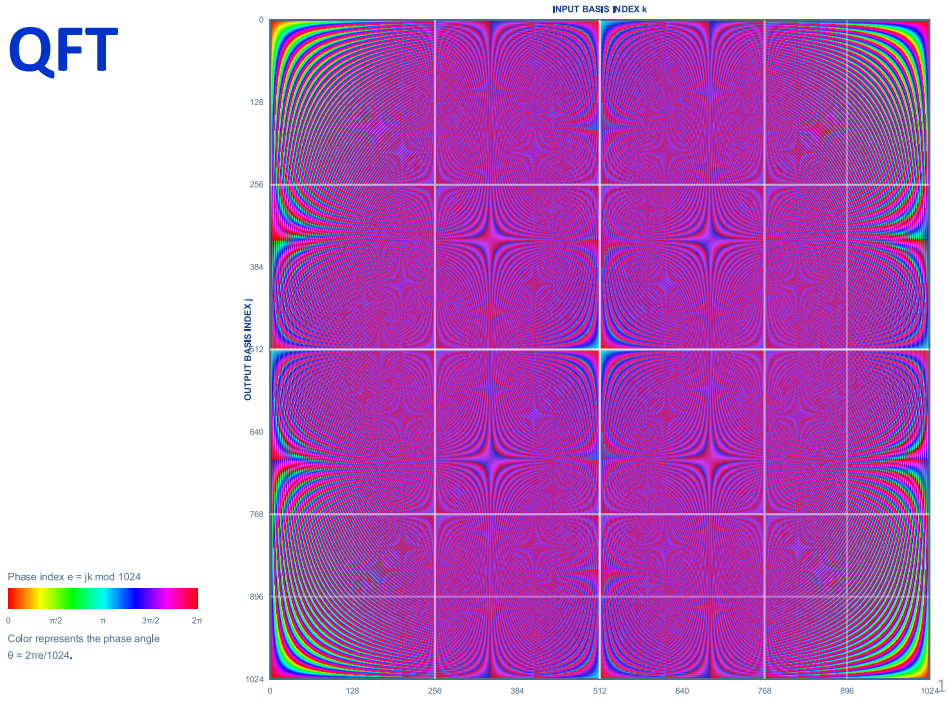

13

For ten qubits, the QFT unitary contains more than one million coefficients. The phase patterns become increasingly intricate. These visualizations make the structure of interference much more tangible.

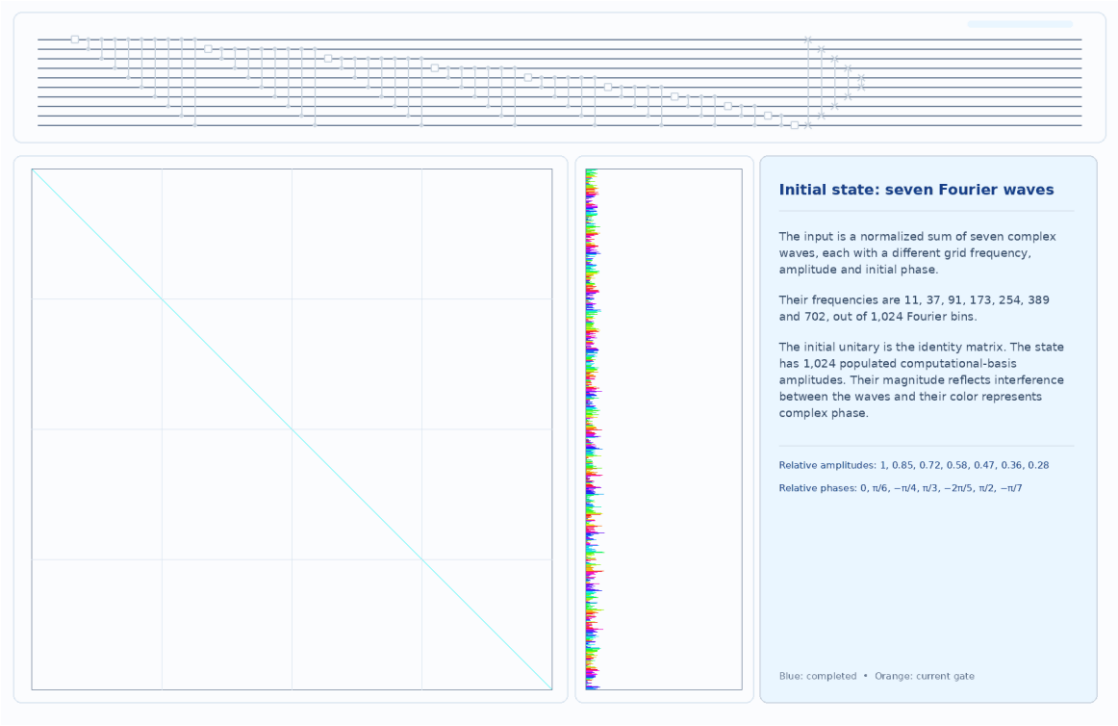

I then asked ChatGPT to animate the successive gate cycles of the circuit. The animation displays both the cumulative unitary and the evolving state vector for an input consisting of seven frequency components. It helps visualize how the desired output emerges as the circuit progresses.

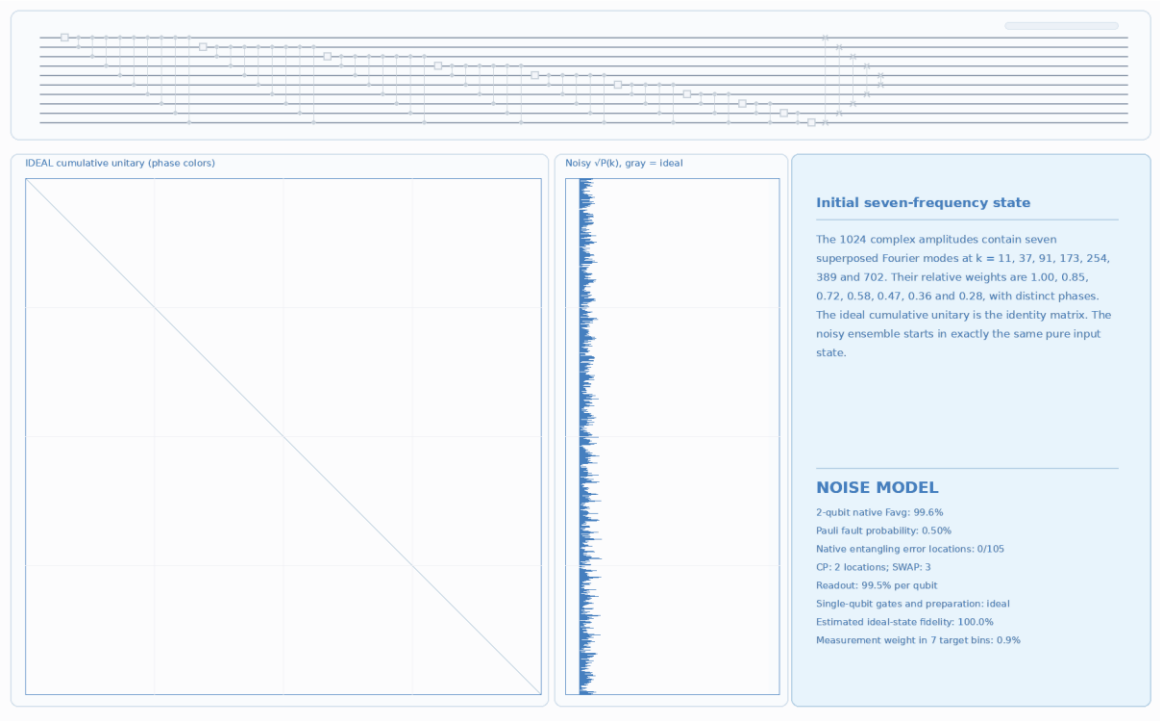

I repeated the exercise with using a qubit noise model. As the gates are applied, errors progressively affect the output distribution. The simulation makes it possible to compare the ideal evolution with the noisy one.

how to get a QFT output result?

you need to “sample” the circuit multiple times

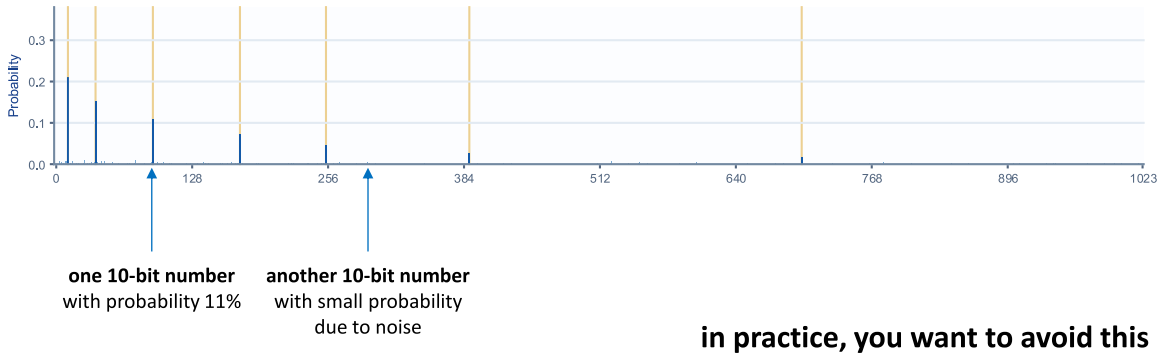

one 10-bit number with probability 11%

another 10-bit number with small probability due to noise

in practice, you want to avoid this

At the end, we do not retrieve the full state vector. Each measurement returns a ten-bit number. A desired frequency may have a probability of 11%, while erroneous results may appear with smaller probabilities. We need to sample the circuit repeatedly to distinguish the useful outcomes from noise. Some algorithms require few repetitions, while a QPE application with poor initial-state overlap or demanding precision can require thousands. The number of shots is an additional resource on top of the cost of each circuit.

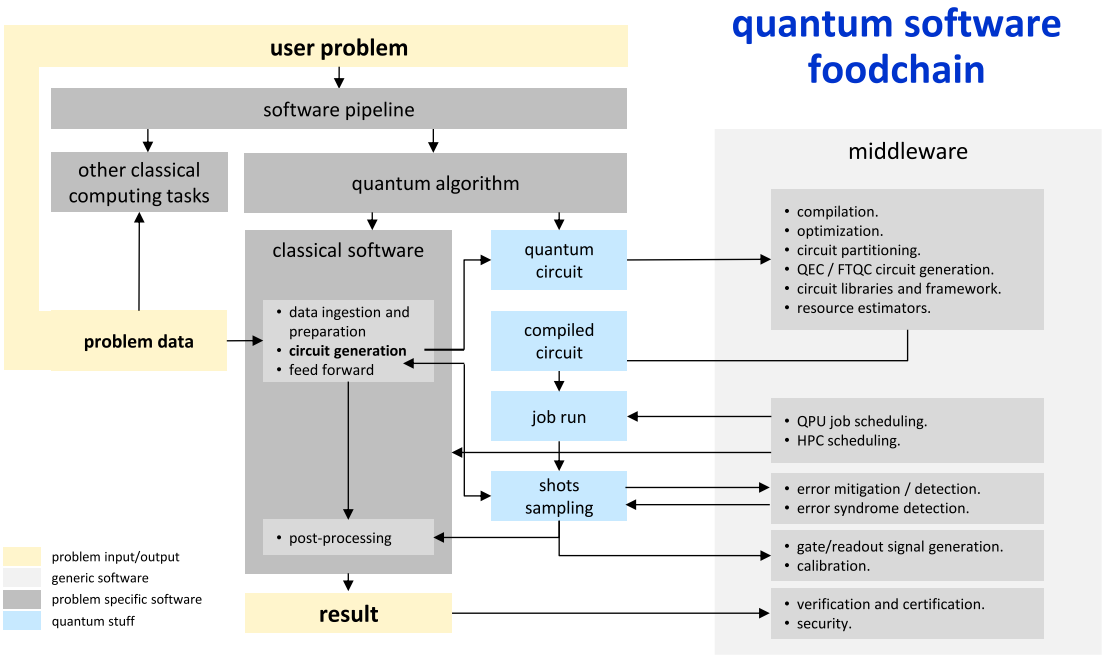

Consider the software value chain for a pharmaceutical company developing a drug. It would not simply run one quantum circuit. A complete pipeline would involve classical calculations and potentially several quantum algorithms, each with preprocessing, circuit generation, compilation, execution, sampling and post-processing.

One important difference with conventional software is the relationship between data and code. We do not recompile an email application each time a new email arrives. In many quantum algorithms, however, the problem data influences the circuit itself. In QPE, changing the simulated many body system may change the state preparation and Hamiltonian encoding, while only the inverse QFT can remain unchanged and in that case, it’s a rather small part of the whole circuit. Parameterized circuits and reusable subcircuits may reduce

the need for recompilation, but the extent of this reuse depends on the application. Compilation can therefore become a variable cost against the user data rather than just a fixed development cost.

The middleware includes compilers, optimizers, circuit partitioning tools, error correction and fault-tolerance software, circuit libraries, resource estimators, job scheduling, calibration, verification, certification and security. Classical processing also takes place during circuit execution, notably for conditional operations and syndrome decoding. This entire software discipline is still much less mature than its classical counterpart.

quantum advantage challenges

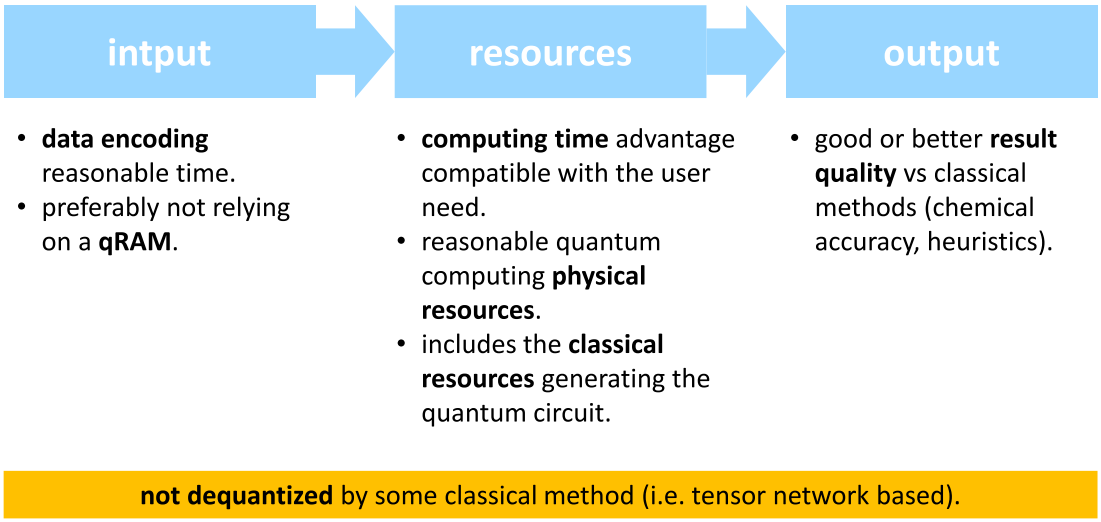

There are several conditions for achieving practical quantum advantage. The first is efficient data encoding. Large input data sets can generate extremely costly state preparation circuits. This is one reason why I am skeptical about running large language models on quantum computers. Their training data can be enormous, and even an already trained model has a very large number of parameters. Efficiently accessing those parameters is a major unresolved problem.

Many theoretical algorithms also assume access to quantum random access memory (qRAM) or an efficient oracle. These assumptions can conceal significant hardware and circuit costs. Chemistry is often more attractive because its useful input description can be relatively compact compared with the size of the quantum state being simulated.

The computing time must also meet the user's requirements. An electricity utility may need a grid optimization result within minutes, whereas a pharmaceutical researcher may accept a calculation lasting several weeks if the result is valuable enough. Physical resources, energy consumption and classical computing costs all need to be accounted for. The output must match or improve on the quality of classical solutions, preferably at a competitive cost.

Finally, the classical baseline is moving. Dequantization describes the development of classical algorithms inspired by quantum algorithms that eliminate or reduce a claimed quantum advantage. Ewin Tang's work on quantum-inspired recommendation systems, following the quantum algorithm of Iordanis Kerenidis and Anupam Prakash, is a well-known example. Tensor networks have also considerably improved classical simulation methods.

The FeMoco complex, associated with biological nitrogen fixation, has been a benchmark in quantum chemistry for many years. Published quantum resource estimates have involved roughly 1,500 to 2,000 logical qubits and very large non-Clifford gate counts, depending on the model and accuracy. Recent classical work discussed in the lecture is challenging the usefulness of this benchmark as an example of quantum advantage. These comparisons depend critically on the active space, target accuracy and computational task.

The implication is that quantum computing must keep moving toward harder problems rather than relying on fixed historical benchmarks.

FTQC resource estimates

FTQC resource estimations

The Grand Challenge of Quantum Applications

Ryan Babbush,* Robbie King,† Sergio Boixo, William Huggins, Tanuj Khattar, Guang Hao Low, Jarrod R. McClean, Thomas O'Brien, and Nicholas C. Rubin
Google Quantum AI, Santa Barbara, CA 93111, United States
(Dated: December 5, 2025)

arxiv.org/abs/2511.09124

Horizon	Problem	Stage III applications	Problem size	Stage IV(b) Logical q.	Stage IV(b) Toffolis	Stage IV(c) Phys. q.	Stage IV(c) Time
Medium	Trapdoor claw-free functions [104]	Classically verifiable quantum advantage	1024 bits	1600 [141]	$8 \cdot 10^6$	–	–
	DQI for OPI [58]		4095 variables 70 constraints	1900	$6 \cdot 10^6$	$8 \cdot 10^5$ [142]	1 hour
	Topological data analysis of random graphs [62]	Unknown	256 vertices 16-cliques	500 [143]	$7 \cdot 10^9$	–	–
	Factoring [34]	Breaking RSA	2048 bits	1400	$7 \cdot 10^9$	$1 \cdot 10^6$ [44]	1 week
	Binary discrete log [34]	Breaking outdated cryptosystems	233 bits	3000	$4 \cdot 10^6$	$4 \cdot 10^6$ [36]	8 min
Furthest	Elliptic curve cryptography [34]	Breaking elliptic curve cryptography	256 bits	3000	$5 \cdot 10^7$	$9 \cdot 10^6$ [35]	4 hours
	Tensor PCA [30]	Unknown	100 variables	900 [144]	$1 \cdot 10^{15}$	–	–
	XOR-SAT refutation [82]	Unknown					
	QAOA for max 8-SAT [102]	Unknown	179 variables	28000	$5 \cdot 10^{11}$	$7 \cdot 10^7$ [145]	15 hours

The Google paper The Grand Challenge of Quantum Applications illustrates the range of logical resources required by candidate quantum algorithms. Resource estimates are often expressed using two key metrics: logical qubit count and the number of T or Toffoli gates. The latter are particularly important because non-Clifford operations are expensive in many fault-tolerant architectures.

physics simulation resource estimates

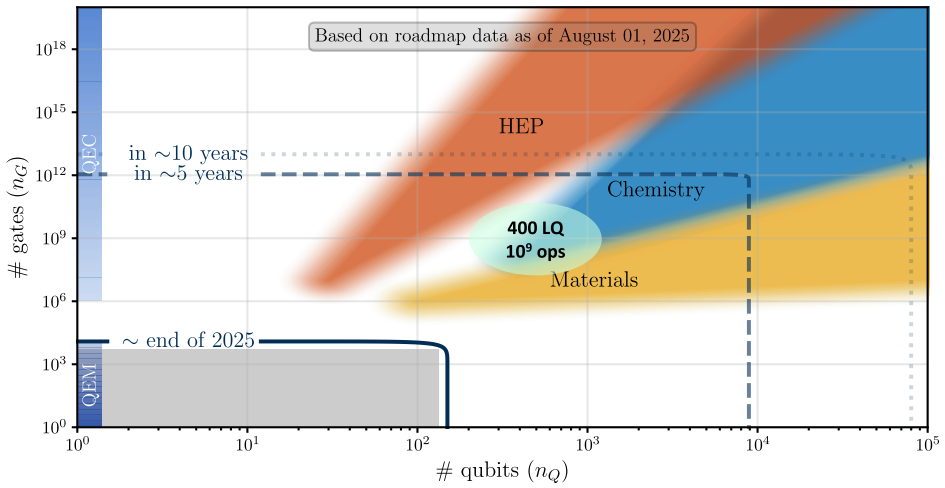

Quantum Computing Technology Roadmaps and Capability Assessment for Scientific Computing -- An analysis of use cases from the NERSC workload by Daan Camps, et al, arXiv, September 2025 (69 pages).

20

The NERSC study of scientific computing workloads covers high-energy physics, quantum chemistry and materials science. The diagram suggests that some potentially interesting workloads begin around a hundred

logical qubits and a million operations, with an illustrative target of approximately 400 logical qubits and a billion operations. These are not universal thresholds, but they provide orders of magnitude.

Present noisy intermediate-scale quantum (NISQ) computers remain far below these targets for reliable general-purpose circuit execution. Error mitigation and error detection can extend the usable circuit size, but there is a large gap between today's demonstrations and the resource requirements of useful fault-tolerant workloads. Both axes of the diagram are logarithmic.

operations fidelity needs

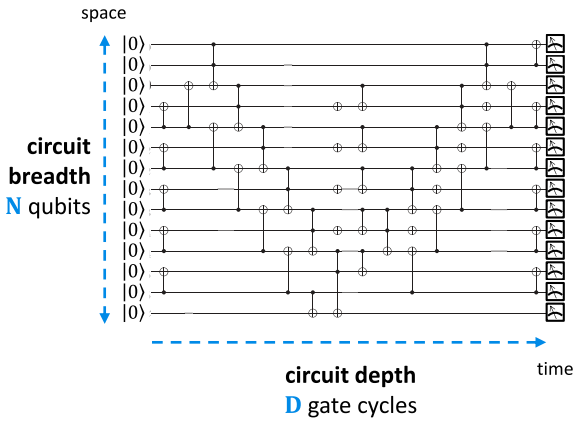

$$\epsilon_{operations} < \frac{1}{N \times D}$$

corresponds to all qubit operations:

- single qubit gates errors
- two-qubit gates errors
- idle state errors
- readout errors

(cc) Olivier Ezratty

The basic fidelity requirement can be illustrated using a circuit with N qubits and depth D. In a simplified independent-error model, the error probability per operation or idle location should be much smaller than 1/(N × D) if we want a high probability of an error-free execution. The empty locations matter because idle qubits are also affected by noise and decoherence.

We need to account for single-qubit and two-qubit gates, idle errors, measurement and reset. Readout and two-qubit gates are often among the most difficult operations, although this varies with the modality. Spatial and temporal variability also matter. A chip's best reported two-qubit gate fidelity is not necessarily representative of the fidelity available across the processor during a long computation. A realistic estimate should consider distributions, correlations and calibration drift rather than only the best or average value.

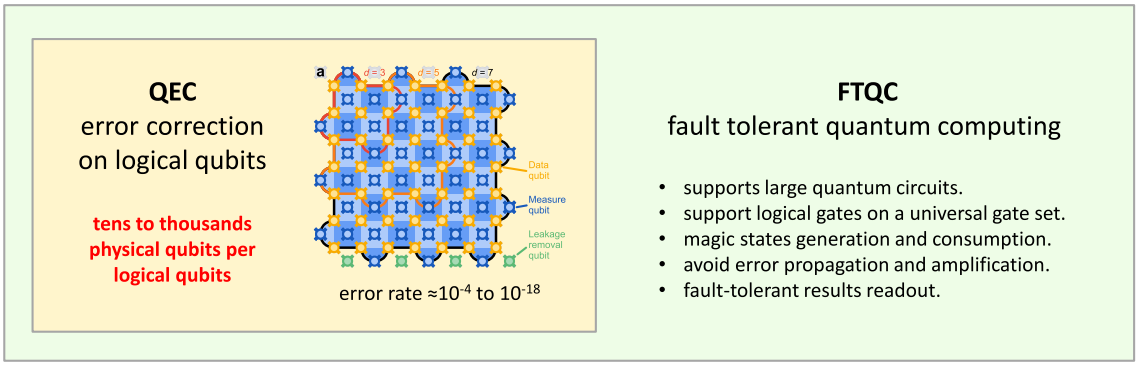

Quantum error correction is not the same thing as full fault tolerance. QEC encodes logical information into multiple physical qubits and protects it against errors. FTQC also requires fault-tolerant initialization, a universal set of logical operations, prevention of uncontrolled error propagation and reliable readout. In many architectures, magic-state generation and consumption are required to implement non-Clifford gates. The number of physical qubits per logical qubit can range from tens to thousands, depending on the code, noise and target logical error rate.

The idea of FTQC was created by Peter Shor back in 1996. Already 30 years!

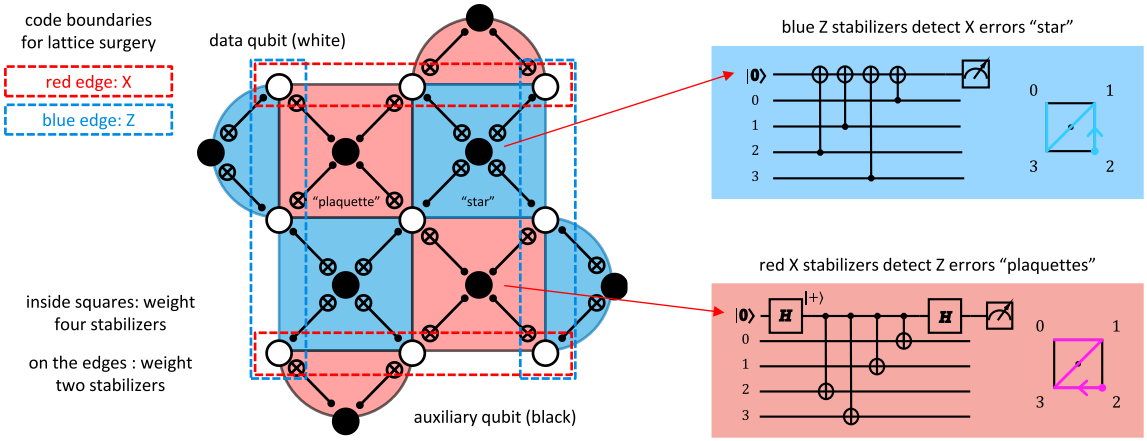

distance 3 rotated surface code with 9 data qubits (white) and 8 auxiliary measurement qubits (black). CNOT gates indicate the local interactions between the data and auxiliary qubits. blue auxiliary qubits detect X error, and red auxiliary qubits detect Z errors. there are weight-four (inside the rectangle) and weight two stabilizers (on the edges of the code). the red squares are "plaquettes" and the blue square are "stars" due to the direction of the CNOTs. d^2 data qubits + d^2-1 auxiliary qubits = $2d^2$-1 total physical qubits.

https://arxiv.org/abs/2411.10406 (Qolab blueprint)
https://arxiv.org/abs/2409.14765v1

A rotated distance-three surface code contains nine data qubits and eight syndrome measurement ancillas. The ancillas interact with neighboring data qubits to measure stabilizers that reveal information about errors without directly measuring the encoded logical state. Stabilizer measurements can establish the entanglement needed for the encoded state. Repeated syndrome extraction and classical decoding help infer errors and track corrections.

color codes and lattice surgery

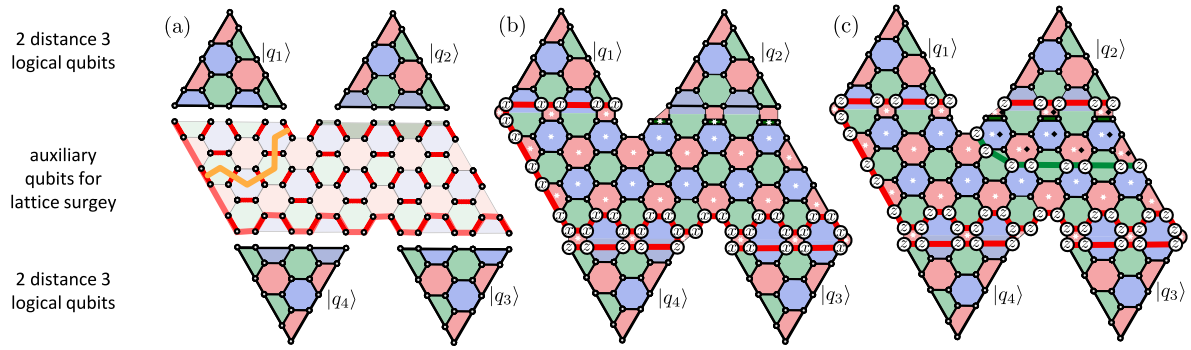

Color codes offer another approach. Some versions have triangular layouts and higher-weight stabilizers than the surface code. They can support transversal Clifford operations such as Hadamard and phase gates more naturally. However, implementing the stabilizers and connecting logical qubits can be more demanding. Lattice surgery uses joint measurements between logical patches, sometimes with additional intermediate patches, to implement logical operations.

pipes representation (Google)

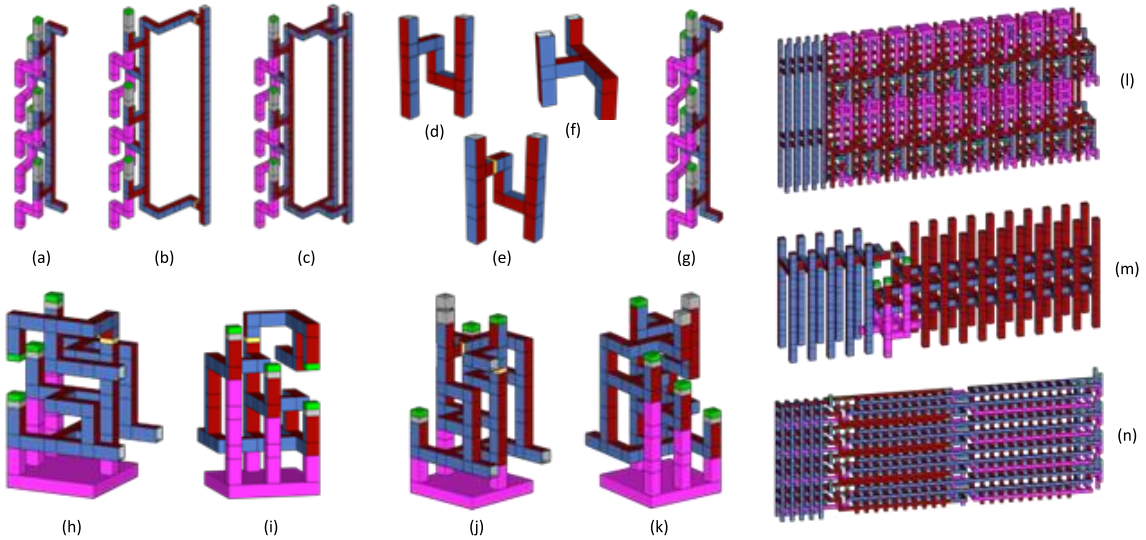

https://arxiv.org/abs/2511.08508

25

Pipe diagrams represent these operations in spacetime. The vertical dimension corresponds to time, and the connections represent interactions between logical patches. These diagrams help describe the organization of logical operations and the additional spacetime volume needed to implement them.

cold-atom shuttling

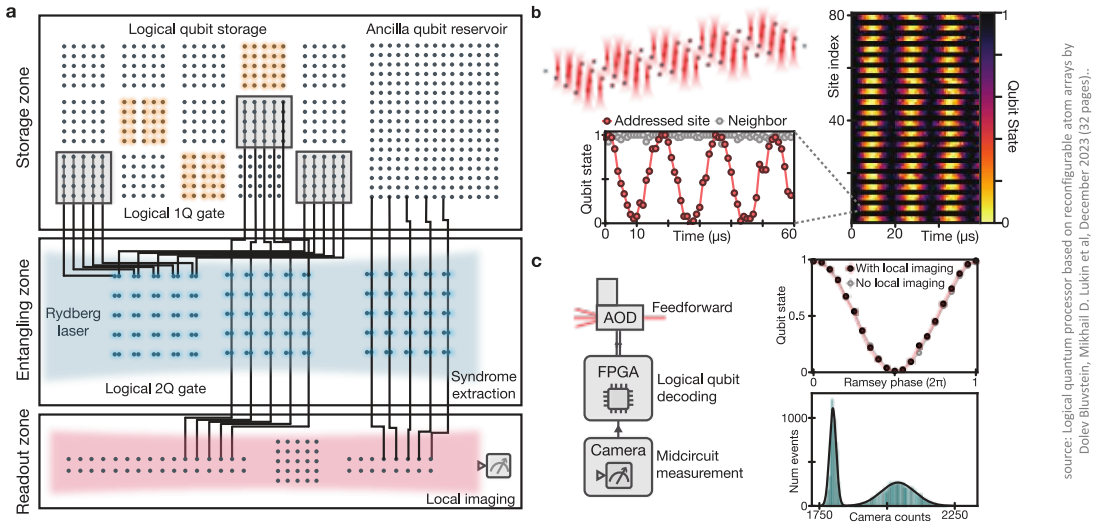

source: Logical quantum processor based on reconfigurable atom arrays by Dolev Bluvstein, Mikhail D. Lukin et al, December 2023 (32 pages)..

26

Neutral-atom systems offer a different approach to connectivity. The December 2023 Harvard and QuEra experiment led by Dolev Bluvstein and Mikhail Lukin demonstrated reconfigurable atom arrays with separate storage, entangling and readout zones. Optical tweezers move selected atoms into an interaction region where a global laser beam can drive Rydberg-mediated entangling gates on multiple pairs. Atoms can subsequently be moved back or to a measurement zone. Replenishment can compensate for atom losses in suitable architectures.

Trapped-ion systems also offer flexible connectivity, sometimes with ion transport. These approaches trade connectivity against the speed of gates, transport and measurement. Solid-state qubits tend to have faster operations but more constrained local connectivity. This is one of many engineering trade-offs.

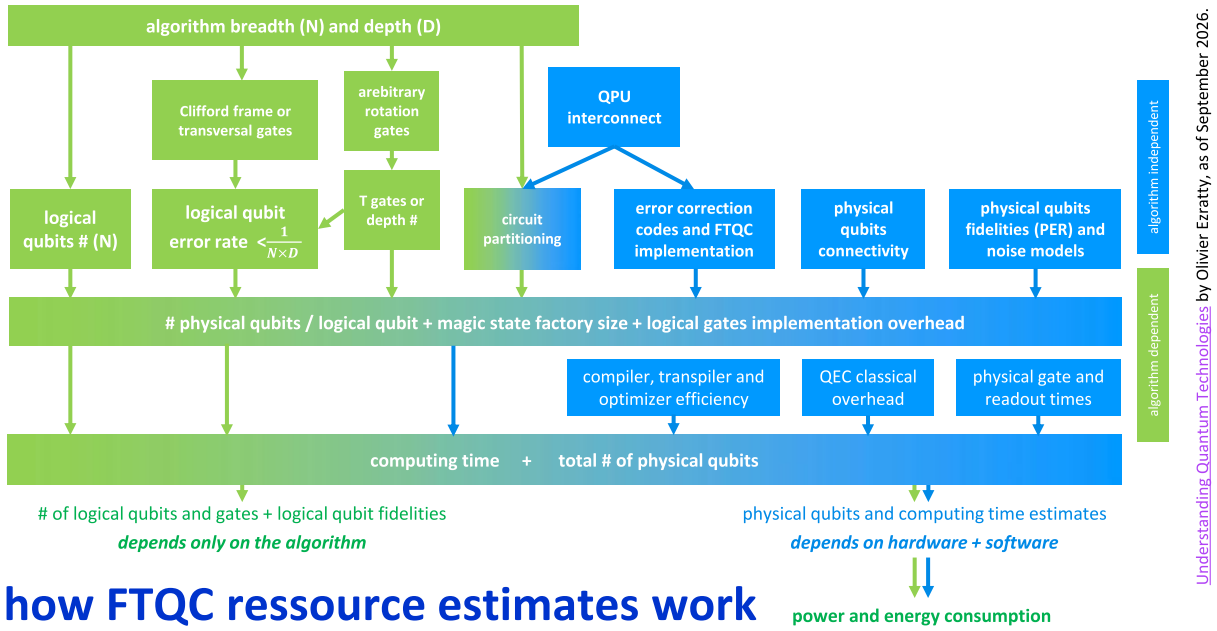

Resource estimation combines application-dependent and architecture-dependent parameters. On the application side, circuit width, depth, the number of rotations to synthesize and the T-depth are important. These determine the required logical error rate and the demand for magic states. On the hardware side, physical error rates, connectivity, error-correcting codes, physical gate speeds, measurement times and interprocessor connections affect the overhead.

The distance of the code must be selected for the logical failure budget and the spacetime volume of the computation. Arbitrary rotations can be synthesized into Clifford and T gates, using methods derived from or improving on Solovay–Kitaev synthesis. The number and scheduling of T gates determine how large the magic-state factories must be. Interconnects and circuit partitioning, such as those investigated by Welinq in France, can add further constraints.

The final outputs are physical qubit requirements and execution time. These can be very different for the same logical algorithm on different architectures. Resource estimation has become a software discipline of its own, with tools developed by companies such as PsiQuantum, Google and Microsoft as well as academic teams.

circuit QEC "shape shifting" example

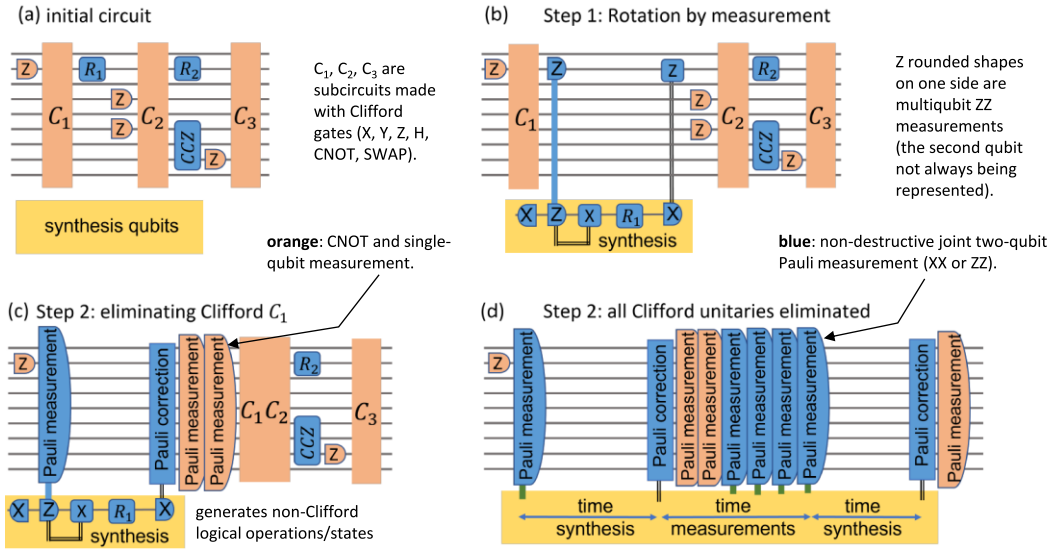

https://arxiv.org/abs/2211.07629

Fault-tolerant compilation can radically transform a circuit. Clifford operations can be commuted or tracked in a Pauli frame, while suitable Pauli measurements and magic-state injections implement the remaining operations. The physical execution is therefore very different from the original logical circuit. The resource estimate must be based on the compiled implementation, not merely the original circuit's gate count. I like to call this process shape-shifting.

QEC/FTQC software + hardware components

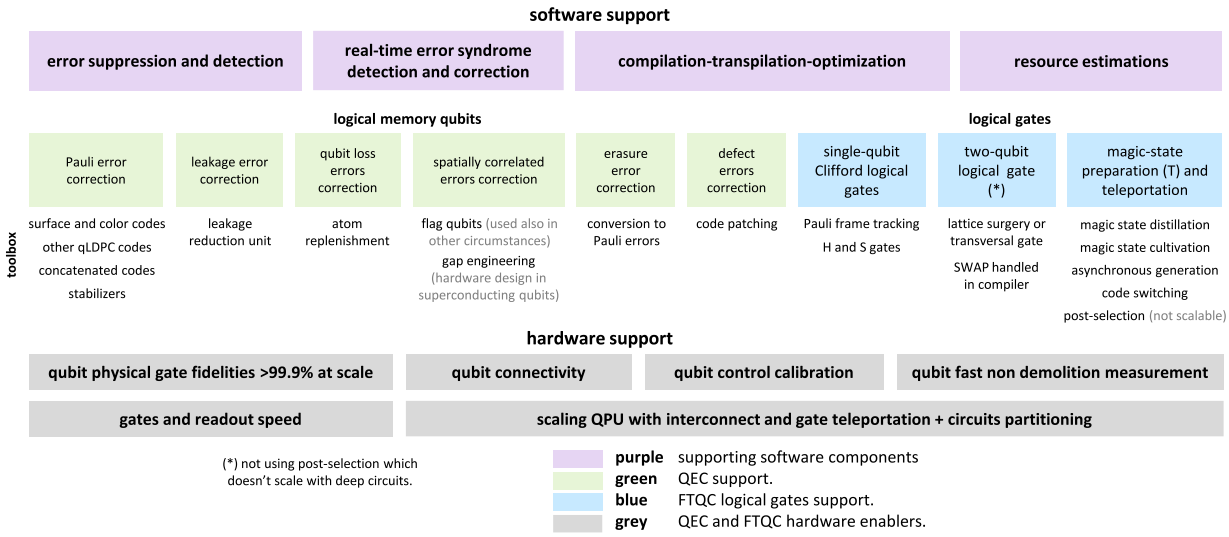

source: Understanding Quantum Technologies 2026, Olivier Ezratty, September 2026

A full FTQC architecture requires logical memory, logical gates, magic-state preparation, connectivity, decoding, compilation and a detailed understanding of the underlying physics. Calibration is particularly important. If a computation lasts days or weeks, we cannot assume that hardware parameters will remain stable throughout. Recent work from Google and IBM on using error syndromes to help adjust calibration during operation is an interesting example of engineering innovation. These problems require close interaction between computer science, hardware engineering and physics.

General information			QEC features						FTQC features											
quantum computing hardware vendor	number of logical qubits	supporting physical qubits	main QEC codes	logical memory experiment	breakeven logical error rate	leakage error correction	erasure error correction	correlated error correction	< breakeven fidelities	single qubit logical gate	singe 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 [1132]	448 [1132]	color codes	[1853]	[1853]	[1132]	[1132]	[1132]	[1853]	[1132, 1853]	[1132]	[1132, 1853]	[1132, 1853]	[1132, 1868]	[1132]	[1132]	[1132]	[1132]	[1132]	[1853]
Infleqtion	12 [1857]	114 [1857]	qLDPC variants	[1857]	[1857]	[1857]	[1857]			[1857]		[1857]	[1859]		CCZ [1859]	CCZ [1859]		[1859]		[1859]
Atom Computing	24 [1869]	208 [1869]	Bacon Shor, 4D code, and other codes [1869]	[1869]			[1869, 1870]			[1869]		[1869]						[1869]		
Pasqal	2	4																yes		
Quantinuum	48 [1871]	98 [1397]	sc, Tesseract, Iceberg, concat, C_4-Helix	[1697, 1855, 1872]	[1872, 1873]		[1855]		[1872, 1873]	[1855, 1856]	[1873]	[1855, 1856, 1872]	[1872, 1873]		[1858]	[1858]		[1855, 1872]	[1872, 1873]	[1858]
IonQ	12 (2026)	256 (2026)	qLDPC bb [1874], walking cats [1826]	[1875, 1876]	[1875, 1876]		[1877]		[1876]		[1826]		[1826]		[1826]	[1826]				[1826]
AQT	2	30	cc, holographic codes [1878]	[1879, 1880]	[1881]	[1882]				[1880, 1883]	[1880]	[1847, 1883, 1884]	[1885]	[1881, 1883]	[1847]	[1847]	[1847]			[1847]
Google	1 [1642]	105 [1642]	sc, cc	[1642]	[1642]	[1642]		[1642]	[1642]					[1886]	[1886]		[1886]	no		
IBM	200 (2029)	n×120	gross code (bb)		[1613]	[1613]				[1613, 1887]		[1613, 1887]		[1613]	[1613, 1888–1891]	[1613]	[1892]			[1613, 1891, 1892]
IQM	4-36 (2026)	54 (2025)	bb, tile codes (qLDPC)	[1615]						[1893]		[1893]								
Alice & Bob	100 (2030))	2000 (2030)	LDPC, elevator codes							[1894]		[1894, 1895]		[1896]	[1896]					
PsiQuantum	100 (2030))	1M (2030)	qLDPC, sc, dual rail				[1897, 1898]			[1897, 1898]		[1897, 1898]		[1899]	[1897–1899]					
Xanadu			GKP, qLDPC, bb				[1900–1902]			[1900–1902]		[1900–1902]			[1900–1902]					
Quandela		12	Floquet							[1903]		[1904]								
ORCA Quantum		12	qLDPC, sc			not applic	[1905]			[1905]		[1905]								

post-selection

source: Understanding Quantum Technologies 2026, Olivier Ezratty, September 2026

30

My inventory of experimental QEC and FTQC capabilities compares the different building blocks demonstrated by vendors. QuEra, Quantinuum and AQT are among the companies that have demonstrated a broad range of capabilities in this inventory. However, the combination of all the required operations in a large, universal fault-tolerant circuit remains the ultimate challenge.

Several experiments use post-selection: they detect errors and discard unsuccessful runs. This is useful experimentally but does not scale well as the circuit size increases. If the acceptance probability decreases exponentially, the number of repetitions increases correspondingly. Nondestructive measurement and active error correction are therefore important milestones.

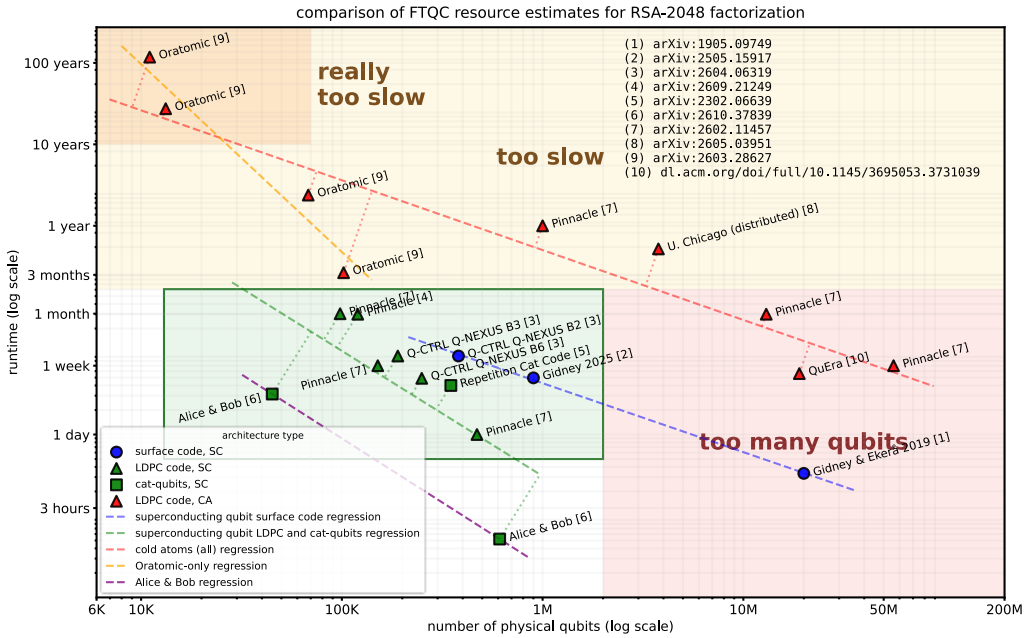

31

I have been tracking resource estimates for RSA factoring, plotting physical qubit count against computation time. These estimates often show a trade-off between space and time. One particularly striking example is an Oratomic estimate using around 10,000 atoms but requiring 118 years of computation. The small qubit

count attracts attention, but the computation time is critical. A century-long computation is not a practical cryptanalytic capability.

There are also engineering limitations unrelated to qubit coherence. The mean time between failures (MTBF) of lasers, electronics and other components must be taken into account for long computations. A system with many components is likely to experience hardware failures during a very long run unless it is designed to tolerate or recover from them.

The more interesting region of the chart involves computation times of weeks or months and manageable physical qubit counts. A recent Alice & Bob estimate discussed in the talk suggests factoring RSA with 32,000 physical qubits in approximately three days, but it assumes cat-qubit noise performance substantially better than current demonstrated devices. Such estimates are conditional designs rather than demonstrations of existing capability.

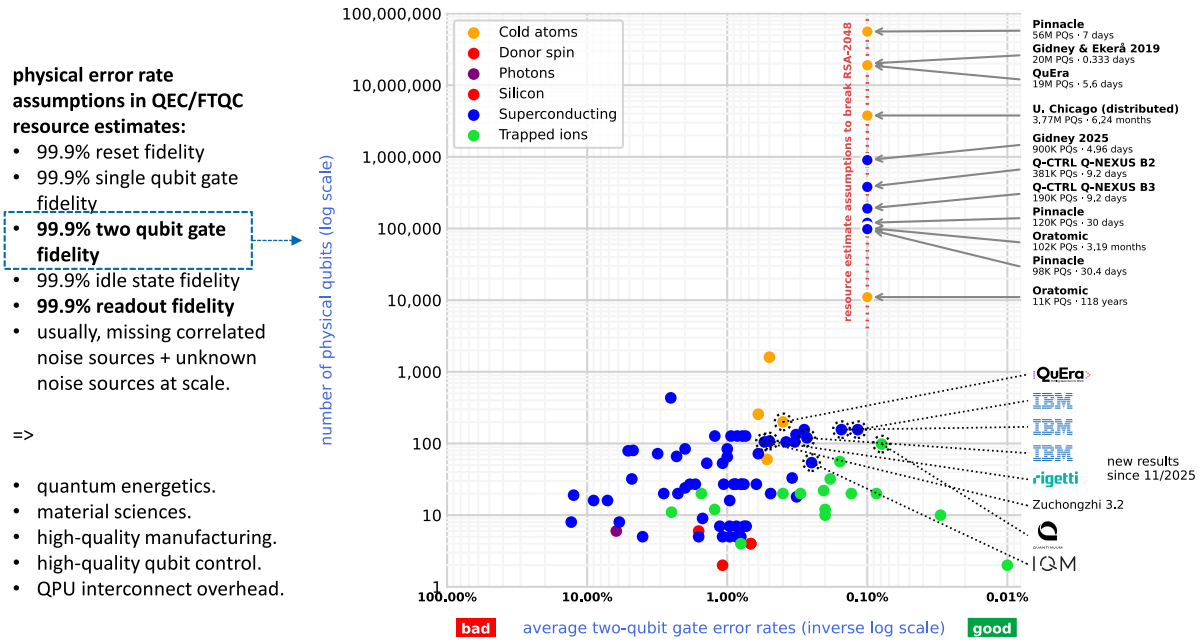

32

The comparison with experimental two-qubit gate fidelities reveals a large gap. The latest results in my inventory are concentrated around tens to roughly a hundred physical qubits, with fidelities ranging approximately from 99.2% to 99.9%. Many resource estimates assume 99.9% fidelity for all relevant operations, including initialization, gates, idle periods and readout. They often omit correlated errors and noise mechanisms that may emerge at scale. Scaling to very large processors while maintaining those assumptions is a major scientific and engineering challenge.

QEC / FTQC space and time overheads

	space cost		time cost
QEC codes	QEC encoding rate	qubit connectivity	number of QEC rounds / cycle
error syndrome decoding	from a laptop to an HPC		classical computing cost
two-qubit logical gates	surgery (2D)	qubit connectivity	atom shuttling (atoms, ions)
magic states generation	T state factory size		can be asynchrone
magic state consumption	gate teleportation (2D)	qubit connectivity	transversal logical gates
physical gate speed			≈ 10 ms to 100 ms (ions or atoms shuttling) / ≈ 10 ns to 300 ns (superconducting qubits)
readout speed	classical readout control chain		variable
mitigating physical noise	QPU interconnect		QPU interconnect latency, LR bound

source: Olivier Ezratty, October 2026

FTQC architectures involve numerous space–time trade-offs. Lattice surgery and gate teleportation consume physical space and syndrome cycles. Transversal logical gates can reduce some overheads, but flexible atomic connectivity may involve slower gates and transport.

Magic-state factories provide another example. More factories can deliver states faster and reduce the time the logical circuit waits for them, at the expense of additional physical qubits. Fewer factories reduce the hardware footprint but increase waiting time. Longer execution can require stronger error correction, offsetting part of the savings. The optimal solution therefore requires considering the system as a whole rather than optimizing individual components independently.

Hardware roadmaps and challenges

FTQC hardware roadmaps and challenges



quantum & classical computing paradigms

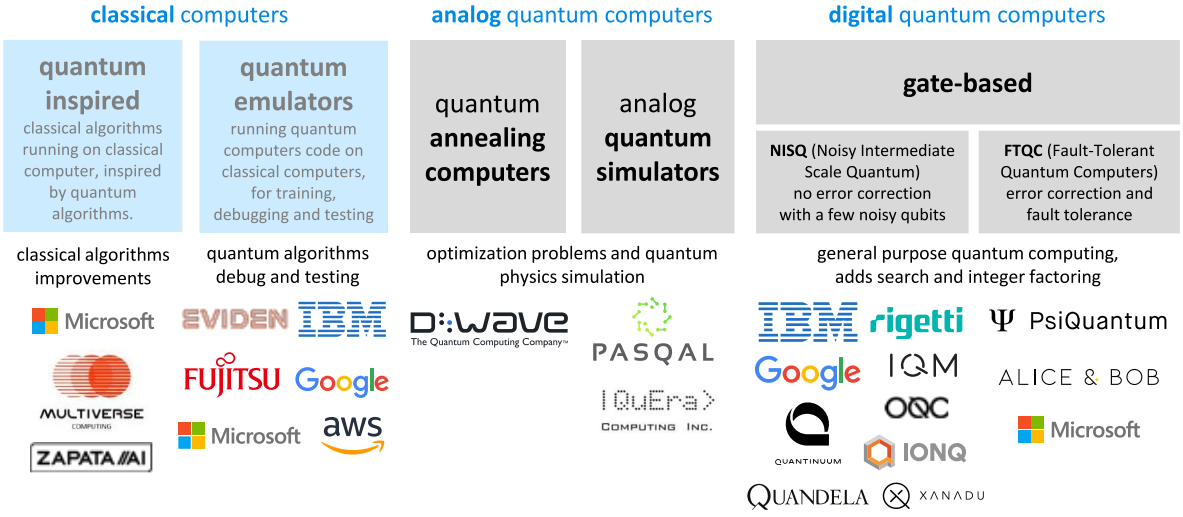

(cc) Olivier Ezratty, 2023

The market includes noisy gate-based computers, fault-tolerant architectures under development, analog quantum simulators, annealers, classical emulators and quantum-inspired classical techniques. These are different tools, and it is important to distinguish them when assessing claims of computational performance.

from NISQ to FTQC

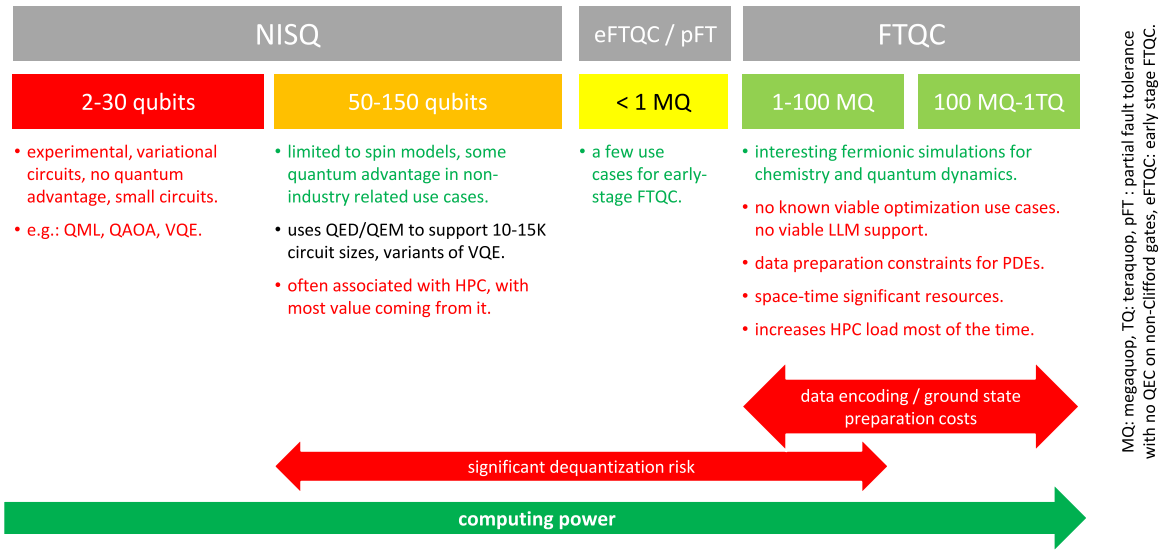

36

NISQ machines can already contribute to scientific experiments, notably simulations of certain spin models. However, much of the algorithmic research on small circuits is primarily useful for learning and experimentation rather than industrial advantage. Many of these circuits can be simulated on conventional computers.

Early FTQC will initially support limited numbers of logical operations. The progression from megaquops to gigaquops and eventually teraquops would open up increasingly complex workloads. Chemistry remains among the most credible application areas, but even there the practical advantages need to be established against strong classical methods.

>90 QPUs industry vendors!

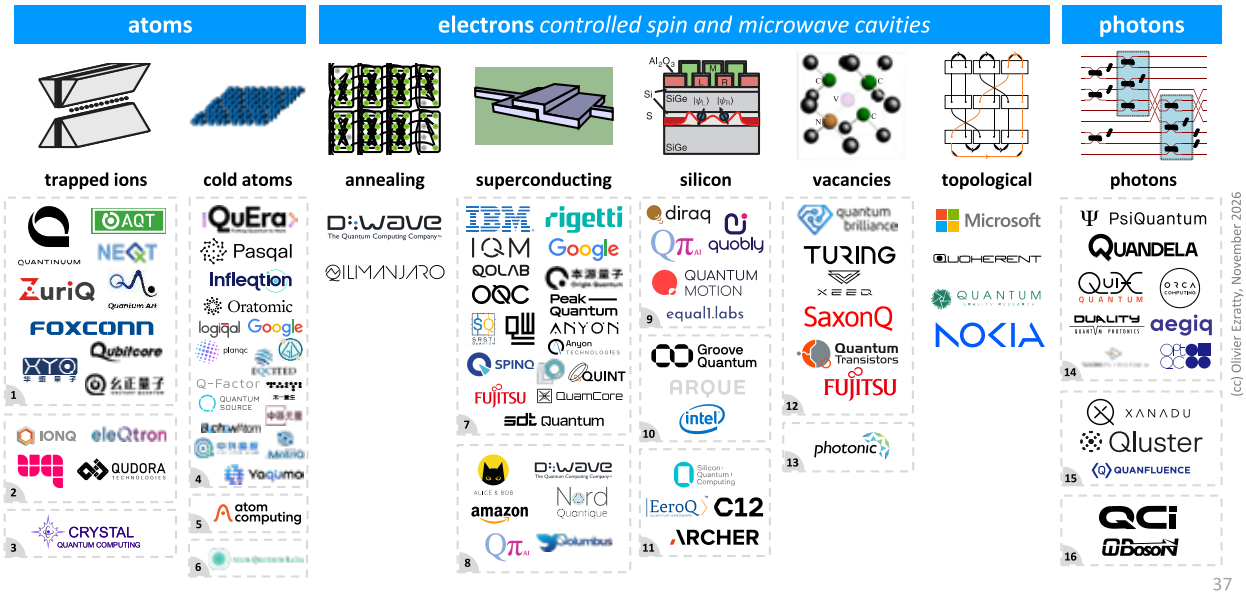

More than 90 companies are developing quantum processors. I asked the audience how many already provide an operational system in the cloud. My own inventory indicates around 20 to 25, a minority of the total. The technological diversity is extraordinary: superconducting circuits, spins, trapped ions, neutral atoms, photons and other modalities rely on very different physical principles and engineering techniques.

I have rarely seen such technological diversity in the IT sector. It makes predicting the winning architecture particularly difficult. Founders naturally believe in their own approach, but from the outside I see significant challenges for all of them. I cannot identify a clear winner at this stage.

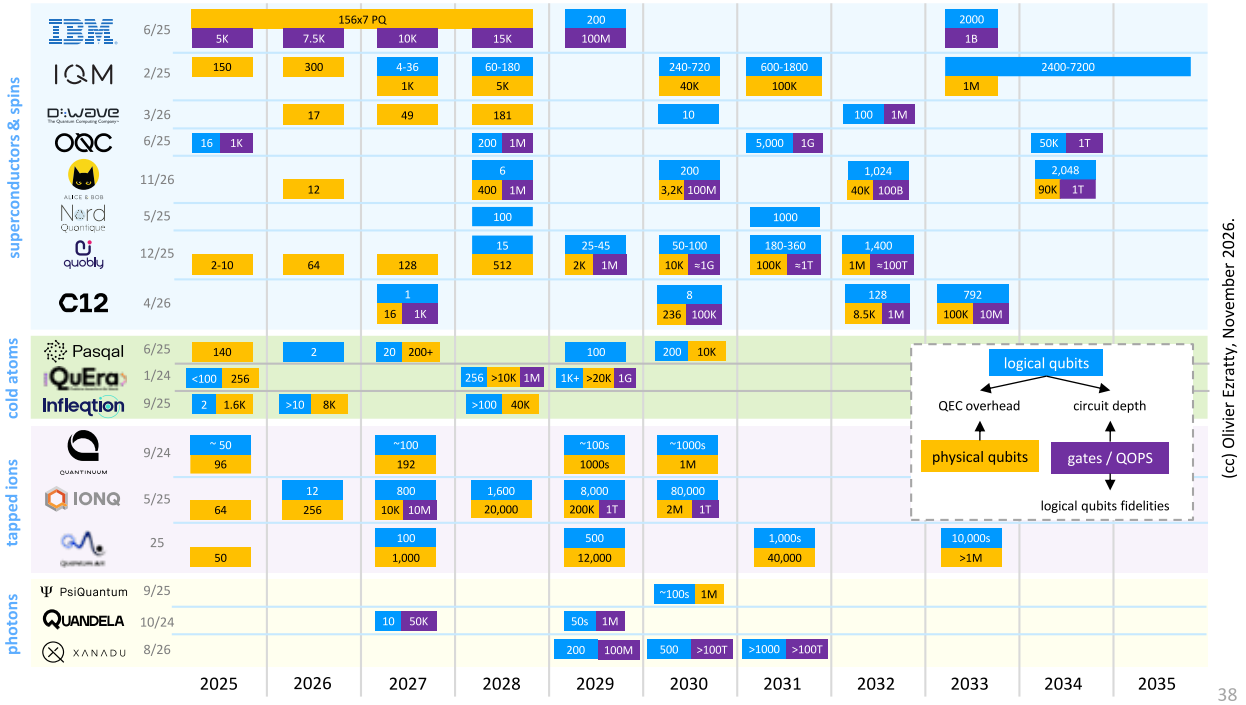

Vendor roadmaps are difficult to compare. The useful metrics include logical qubit count, logical error rates or achievable circuit size, physical qubit count and execution speed. A fixed number of physical qubits can support more logical qubits in a shallow circuit or fewer logical qubits in a deeper circuit if stronger

protection is necessary. A roadmap should therefore specify the workload associated with its logical qubit figures.

These are promises, not measured capabilities. I expect many companies to be late because chip design, manufacturing, characterization and system integration involve repeated iterations. For example, Google's superconducting processors went from approximately 50 qubits in 2019 to 105 in Willow five years later. Very rapid projected increases in physical or logical qubit count deserve scrutiny.

key hardware challenges

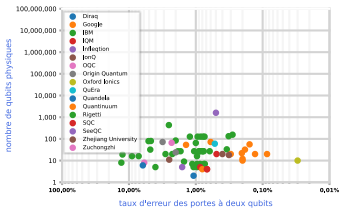

qubit quality

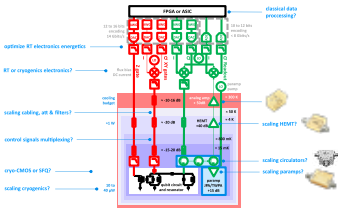

electronics, cabling, cryogeny

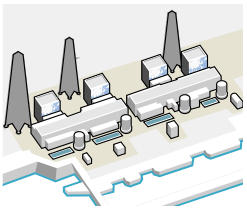

cost and power/energy

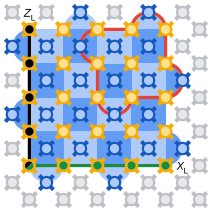

error correction

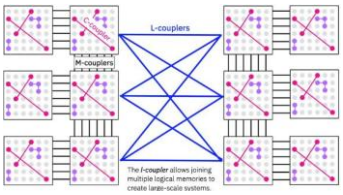

interconnection

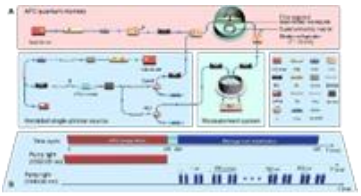

quantum memory

39

The main hardware challenges are improving physical qubit quality, implementing error correction, connecting processors, developing control electronics, cabling, photonics and cryogenics, and controlling power consumption and cost. Quantum memory could also play a useful role in some architectures. Each of these challenges can keep scientific and engineering teams busy for years.

qubit errors sources

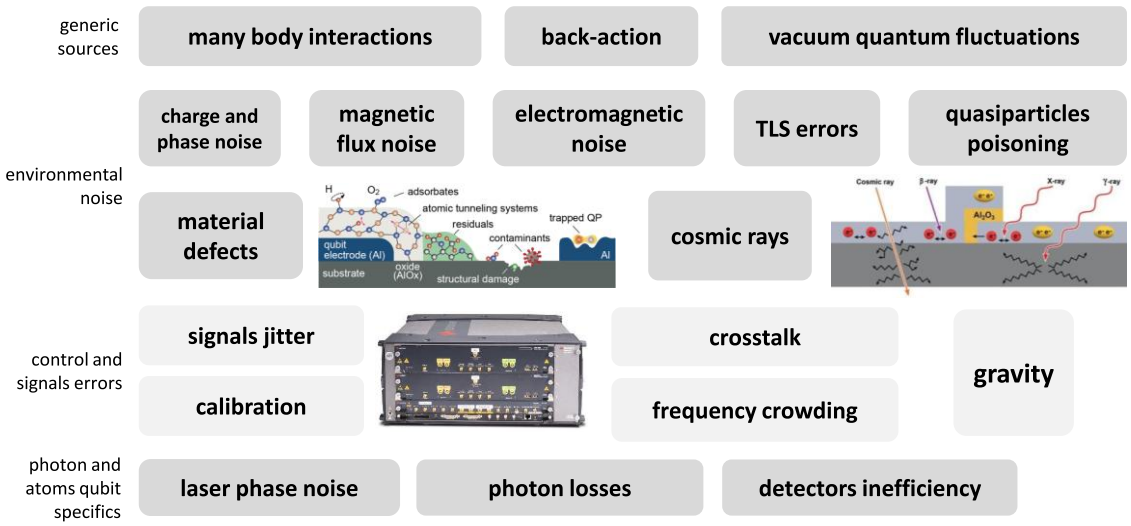

(cc) Olivier Ezratty, 2023-2024

Physical errors have many origins: material defects, electromagnetic noise, charge and flux fluctuations, quasiparticles, crosstalk, calibration drift, laser phase noise, photon losses and detector inefficiencies. They differ across modalities.

An audience question concerned cosmic rays. Their impact is particularly documented in superconducting circuits, where energetic particles can create quasiparticles and correlated errors. Other solid-state devices may also be affected. I am not aware of comparable evidence of a major direct cosmic-ray limitation in neutral-atom or trapped-ion systems. However, those platforms have other sources of correlated noise. A global laser beam, for example, can simultaneously affect all atom pairs addressed by that beam if its phase or intensity fluctuates.

Quantum computing energetics

quantum computing energetics



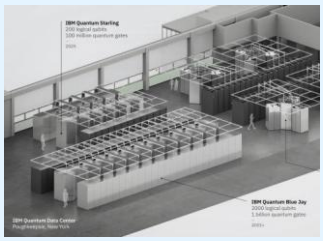



#QEI
the quantum energy initiative

quantum-energy-initiative.org/

why

- ICT energy consumption is growing fast.
- most quantum technologies will add an incremental energetic cost to existing ICT.
- it is time to work on this as quantum technologies and applications are designed.
- responsible innovation imperative.





QEI 2025 (Grenoble) and 2026 (Barcelona)

what

- build a **new** science and engineering field.
- create full-stack methodologies and standards to **evaluate**, **optimize**, and **benchmark** quantum technologies energy consumption.
- favor integrative **full-stack approaches**.

how

- QEI scientific board (10).
- QEI community (545 participants from 67 countries) *.
- QEI workshops (2023 in Singapore, 2025 in Grenoble, 2026 in Barcelona, 2027 in Lisbon).
- QEI IEEE standards Workgroup.
- ≈30 industry partners including IBM, Quandela, etc.

Energetics accounts for a large part of my current activity. We created the Quantum Energy Initiative four years ago to combine fundamental physics with a full-stack engineering approach. The idea is to investigate energy consumption while quantum computers are being designed, rather than discovering the problem once systems consume enormous amounts of power, as has happened to some extent in AI.

Our objective is to build an interdisciplinary field capable of evaluating, optimizing and benchmarking quantum technologies' energy consumption. We connect academic and industrial teams and contribute to standardization efforts defining the scope of energy accounting. The community meets regularly, including at the annual QEI workshop. The most recent workshop was in Barcelona, and the next is planned in Lisbon.

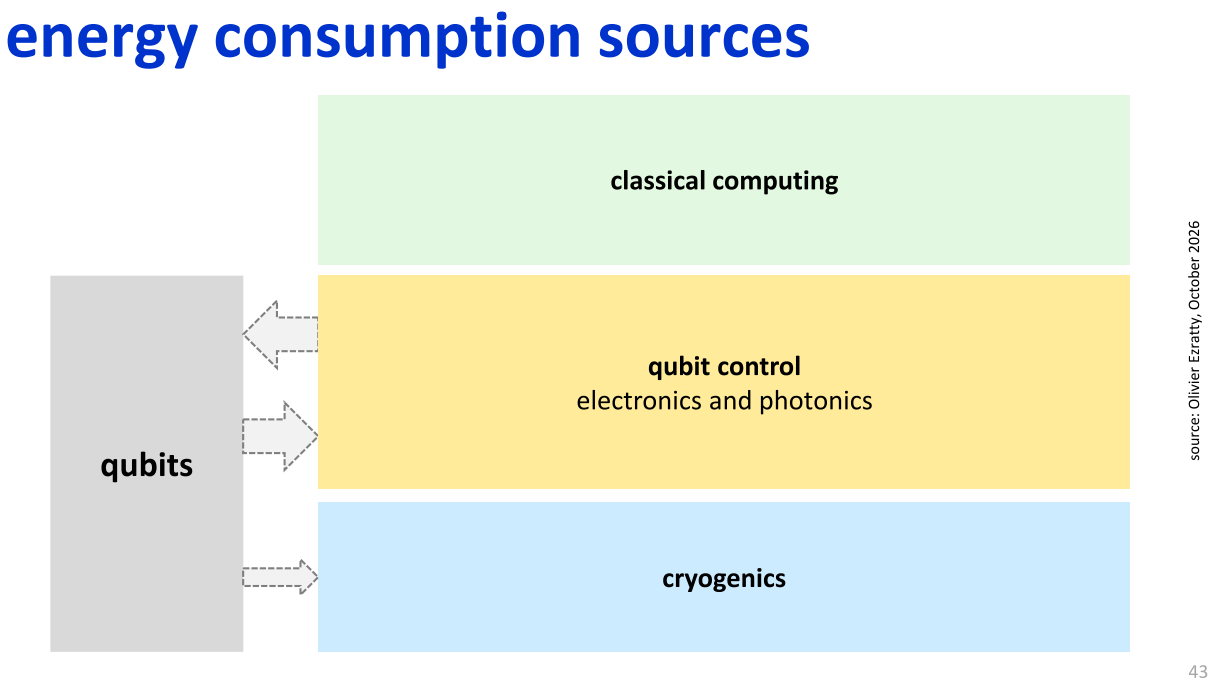

Most power is consumed by macroscopic systems that control, cool and process information associated with qubits. Three main categories are cryogenics, control electronics and photonics, and classical computing. The latter includes both application-related computation and the computers controlling the quantum hardware. The key question is how these costs evolve when we scale toward millions of physical qubits.

quantum ↔ macroscopic costs

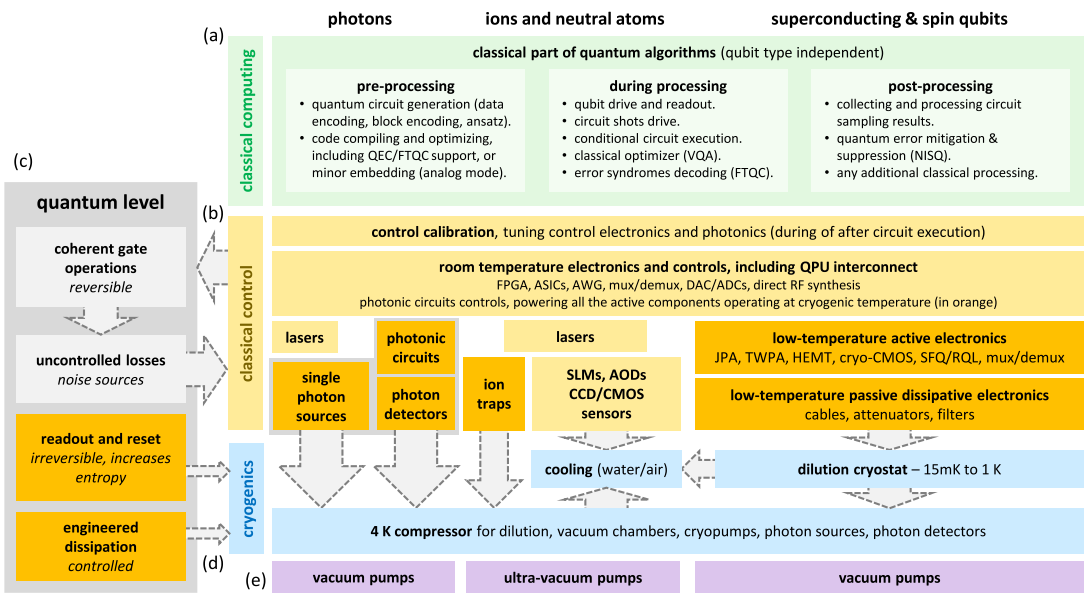

source: Understanding Quantum Technologies 2026, Olivier Ezratty, September 2026

44

The detailed architecture differs between solid-state, atomic and photonic qubits. Solid-state platforms require room-temperature electronics as well as passive and potentially active cryogenic components. Atomic systems require lasers, beam-steering devices, cameras, vacuum pumps and associated cooling. Photonic systems may require cryogenic single-photon sources or detectors as well as room-temperature photonic controls.

Cooling must be considered across the complete chain, including the cooling of room-temperature lasers, compressors and water circuits. Some neutral-atom systems use cryogenic vacuum chambers to improve vacuum quality. The atoms themselves are cooled with laser techniques rather than by direct contact with the cryogenic surfaces.

At the quantum level, coherent unitary operations are reversible in principle, while measurement and reset involve irreversible processes and entropy management. Engineered dissipation, such as that used to stabilize cat-qubit states, can be a useful resource. Noise and the control systems needed to counter it connect the microscopic physics to macroscopic power consumption.

quantum energetics scientific fields

quantum technologies energetics

full-stack approaches to model, optimize, and benchmark the energy resources required to implement useful quantum technologies, under specified performance constraints.

macroscopic scale	classical computing software engineering quantum information science all classical enabling technologies classical thermodynamics
quantum scale	quantum electrodynamics, semiconductor spin physics, quantum optics. open quantum system theory, quantum control theory, noise modelling. materials science and semiconductor manufacturing. fundamental quantum energetics. quantum thermodynamics.

The disciplines involved range from quantum optics, quantum electrodynamics, spin physics and open quantum systems to control theory, materials science and manufacturing. Materials are particularly important because defects in superconducting circuits and interfaces can limit qubit fidelity. Improving those materials may substantially reduce the overhead of error correction.

Quantum thermodynamics and fundamental quantum energetics investigate the energy associated with quantum processes, including measurement, coherence and entanglement. At the macroscopic scale, classical thermodynamics, electronics, cryogenic engineering, computer science and software engineering all contribute. We need an overarching full-stack approach to identify dependencies and optimize the complete system.

counter intuitive insights



cryogenics

- all qubit types rely on cryogenics.
- it's usually a variable cost vs physical qubit #.
- cryogenics may not be the dominant energetic cost for some qubit modalities.

science vs engineering

- still room for progress at the fundamental level.
- enabling technologies key role.
- many interdependencies and potential counterintuitive optimizations.



electronics and photonics

- control electronics and photonics may become the dominant energetic cost in solid state platforms.

classical computing

- classical energetic costs must be considered before, during and after a quantum circuit run.
- compilation is a variable cost vs the problem data.



error correction

- error syndrome detection may dominate energetic costs for some qubit modalities.

economics

- energy could represent >30% of an FTQC TCO.
- energy = power (infra) × time (user constraints).
- power is a proxy for all environmental footprint.

46

Some of our findings are counterintuitive. Cryogenics attracts a lot of attention, but it is not necessarily the dominant power cost at scale. Low temperatures can improve qubit quality. In semiconductor quantum-dot photon sources, for example, temperature affects photon indistinguishability. Increasing the temperature to save cooling power may therefore increase other costs.

For some projected large superconducting systems, vendor estimates suggest cryogenics could account for less than 20–25% of total consumption, with control electronics becoming a major contributor. Error-syndrome decoding may also be costly, particularly for fast qubits or codes requiring complex decoding. Codes with lower physical-qubit overhead can shift part of the burden to classical computation.

Classical preprocessing, compilation and post-processing must also be included. Energy could become a significant component of total cost of ownership, potentially exceeding 30% in some scenarios. These figures are architecture-dependent estimates rather than universal characteristics of FTQC.

NISQ-FTQC-HPC power comparables

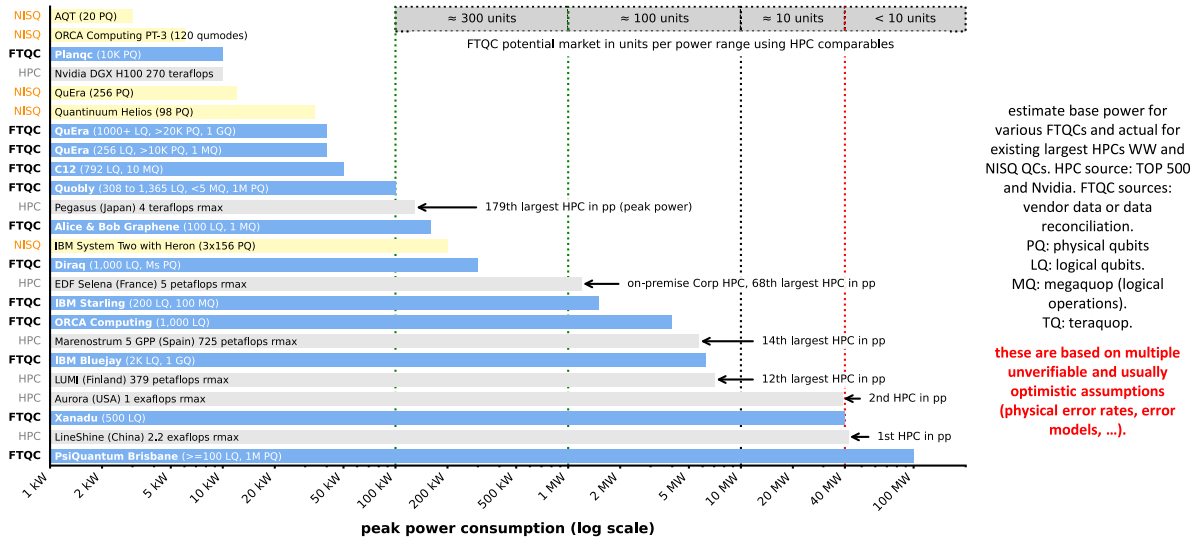

estimate base power for various FTQCs and actual for existing largest HPCs WW and NISQ QCs. HPC source: TOP 500 and Nvidia. FTQC sources: vendor data or data reconciliation.
PQ: physical qubits
LQ: logical qubits.
MQ: megaquop (logical operations).
TQ: teraquop.

these are based on multiple unverifiable and usually optimistic assumptions (physical error rates, error models, …).

47

I maintain a chart comparing the power of current NISQ computers, future FTQC projections and existing high-performance computing (HPC) systems. Current quantum computers generally consume less than 200 kW in the examples shown. Future fault-tolerant systems could range much higher, with some estimates reaching approximately 100 MW. The largest HPC systems already operate at powers of several tens of megawatts. HPC is a particularly relevant comparison because scientific simulation is likely to be one of the major uses of FTQC.

power comparable warnings

1. **vendor numbers give some indication of the potential power/energy needs for future FTQC systems.**
2. **the distribution and comparison with HPC systems gives an indirect indication of a potential market size.**
3. **numbers are based on multiple optimistic assumptions:**
   - circuit sizes and optimistic physical error rates: 99.9% for all operations.
   - incomplete error models: leakage, losses, correlated errors at scale, variability.
   - not accounting for the overhead of quantum interconnect.
   - quantum algorithm success metrics (accuracy, …).
   - enabling technologies efficiencies: cryogenics, electronics, photon losses, QEC.
4. **they do not account for the speed difference between systems.**
   - these differences play a key role when comparing the energy consumption and computing times for specific quantum computing workloads.

These comparisons come with substantial caveats. Vendor projections often lack detailed documentation and use optimistic assumptions concerning physical error rates, error correlations and scalability. Moreover, power is not energy. A low-power system with slow gates can consume more energy per completed computation than a higher-power system that runs much faster. We must also account for classical processing and the complete workload, not just the quantum processor.

why look at both energy and power?

1. **broad range of expected power consumption**
   - 1 to 1000 range under usual optimistic assumptions.
   - maximum = 2,5 the largest exaflops HPCs and 10% of an average nuclear reactor (1 GW).
2. **power consumption**: indication of feasibility.
3. **energy consumption**: potential platform Darwinism through energy-efficiency, « survival of the fittest »?
4. **global environmental footprint**: will depend on # of practical use cases.

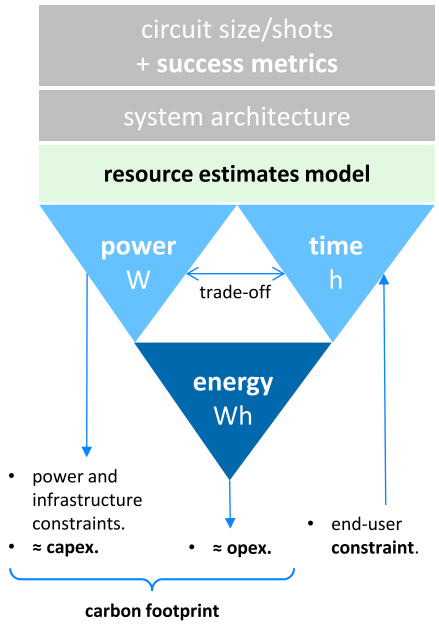

49

Power, energy and time serve different purposes. Power determines infrastructure requirements, including transformers, electrical connections and cooling, and is therefore related to capital expenditure. Energy, the product of power and time for a constant load, contributes to operating expenditure. Time is a user constraint. Parallel execution of independent circuit shots can reduce elapsed time, but may increase the infrastructure needed.

These quantities must be evaluated using comparable workloads and success criteria. A quantum system requiring enormous infrastructure may face economic and deployment limitations even if a theoretical algorithmic advantage exists.

There's also a possibility that physically compact systems with short interconnects may have a scientific advantage through reduced latency, not merely lower cost.

Total sector carbon footprint is unknown and depends on both hardware lifecycle and application-driven market size.

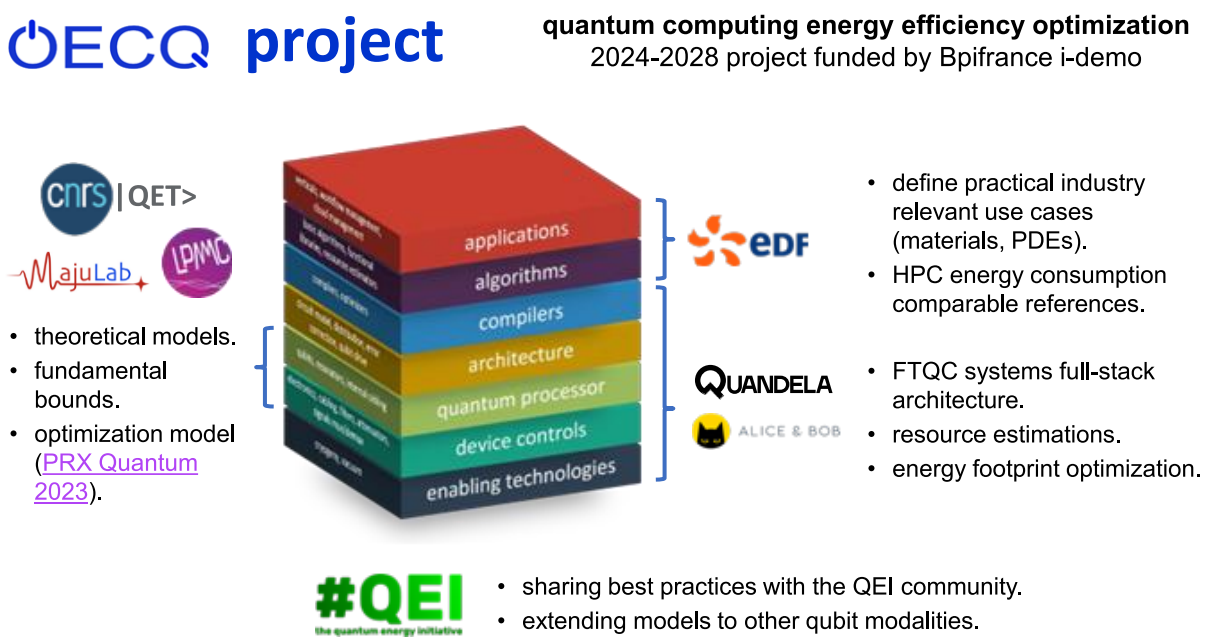

The OECQ project, funded by Bpifrance for 2024–2028, investigates quantum computing energy efficiency. It combines industrially relevant use cases, particularly materials and PDEs, classical HPC baselines, full-stack FTQC architecture models, resource estimates and energy optimization. The aim is to connect fundamental energetic models with practical system engineering and share the resulting methods with the QEI community.

Quantum computing ecosystems



the LLM-based research revolution

QuOptLab: Advancing Quantum Optics Experiments with an AI Agent

Yuqing Li,* Zeguan Wu,† Yu Gan,‡ Min Chen,§ and Junyu Liu¶
University of Pittsburgh

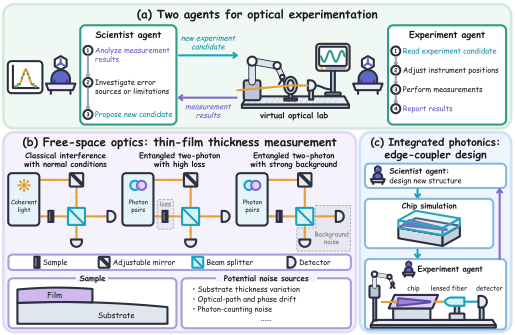

https://arxiv.org/abs/2611.32000

Agentic Pipelines for Mapping Topological Codes to Quantum Hardware

Arpit Arora,[1] Adam Holmes,[1] Krysta M. Svore,[1] and Taylor L. Patti[1]
[1]*NVIDIA, Santa Clara, California 95051, USA*
(Dated: September 1, 2026)

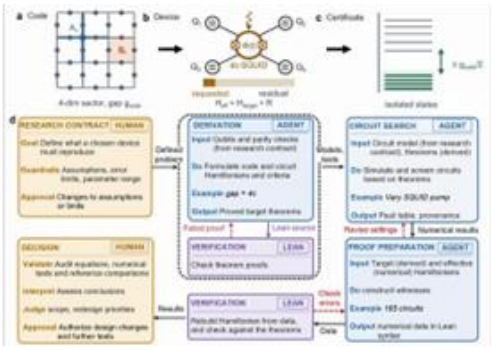

https://arxiv.org/abs/2611.03329

Generative AI is changing research at an extraordinary pace. New papers appear almost every day, reporting AI-assisted mathematical and physical discoveries. This is relevant to quantum computing because it can improve quantum algorithms, error-correcting codes and the classical algorithms with which quantum computers must compete. Two very recent arXiv papers: Q-OptLab, is an AI-agent system for quantum optics experiments from Pittsburgh, and an NVIDIA-related hardware mapping result.

AI Scientist Discovered Robust Quantum Control

Jiahao Yan,[1,*] Xuanying Lai,[1,*] Xiaoshui Lin,[1,*] Yizhou Wang[1], Changyu Yao[1],
Guanghui He[1], Thomas Lewer[1], Chong Zu[1,2,†] Chuanwei Zhang,[1,2,†]
[1]Department of Physics, Washington University, St. Louis, MO 63130, USA
[2]Center for Quantum Leaps, Washington University, St. Louis, MO 63130, USA
*These authors contribute equally to this work
†To whom correspondence should be addressed; E-mail: zu@wustl.edu, chuanwei.zhang@wustl.edu
(Dated: November 18, 2026)

https://arxiv.org/abs/2611.21765

Agentic Pipelines for Mapping Topological Codes to Quantum Hardware

Arpit Arora,[1] Adam Holmes,[1] Krysta M. Svore,[1] and Taylor L. Patti[1]
[1]*NVIDIA, Santa Clara, California 95051, USA*
(Dated: September 1, 2026)

https://arxiv.org/abs/2611.02833

Agentic AI for Discovering Fault-Tolerant Quantum Circuits

Aritra Lahiri[1,*] and Florian Marquardt[1,2,†]
[1]*Max Planck Institute for the Science of Light, 91058 Erlangen, Germany*
[2]*Department of Physics, Friedrich-Alexander University Erlangen-Nürnberg, 91058 Erlangen, Germany*
(Dated: November 16, 2026)

https://arxiv.org/abs/2611.18504

Discovering state-of-the-art quantum circuits with large language models

Andi Gu,[1,2,3,*] Alek H. Kemeny,[4] Susanne F. Yelin,[1,2] and Mikhail D. Lukin[1,2,5]
[1]*Harvard Quantum Initiative for Science and Engineering, Cambridge, MA 02138, USA*
[2]*Department of Physics, Harvard University, Cambridge, MA 02138, USA*
[3]*Anthropic STEM Fellows Program*
[4]*Anthropic*
[5]*QuEra Computing, 1380 Soldiers Field Road, Boston, MA 02135, USA*

https://arxiv.org/abs/2610.22367

Discovering formally verified quantum programs with AI

Sergey Shirobokov[1,*,†], Ben Jaderberg[1,*,†], Katerina Hristova[2,*], George Tsoukalas[1,*], Johannes Bausch[1,*],
Eric Wieser[1], Francisco J.H. Heras[1], Nicholas C. Rubin[2], Guang Hao Low[2], Sébastien Racanière[1],
Nenad Tomašev[1], Lei M. Zhang[1], James D. Watson[2], Tanuj Khattar[2], Aws Albarghouthi[2],
Stavan Jain[2], Harun Khan[1], Leopoldo Sarra[1], Andrew W. Senior[1], Michael Broughton[2],
Hussain Masoom[1], Sergio Boixo[2], Ryan Babbush[2], Swarat Chaudhuri[1], and Pushmeet Kohli[1]
[1]Google DeepMind [2]Google Quantum AI

https://arxiv.org/abs/2611.09822

Reducing RSA-2048 Factoring Resources Below One Billion Toffolis via LLM-Driven Verifiable Quantum Compilation

Xiaodi Li,[1,4,*] Yuqing Wang,[1,2,4,*] Xue Chen,[3,4] Yang-Yang Tang,[3,4] Linghui Chen,[1,4,†] and Tao Zhang[1,3,4,†]
[1]*Intelligent Quantum Inception Co., Ltd., Haidian, Beijing 100085, China*
[2]*Department of Physics and State Key Laboratory of Low Dimensional Quantum Physics, Tsinghua University, Beijing 100084, China*
[3]*Chen-Ning Yang Institute for Advanced Study, Tsinghua University, Beijing 100084, China*
[4]*Beijing Key Laboratory of Cold Atom Quantum Computation, Tsinghua University, Beijing 100084, China*
(Dated: November 22, 2026)

https://arxiv.org/abs/2611.31158

The examples in the slides illustrate the breadth of this research, including mathematical complexity, quantum information and classical simulation. It is increasingly difficult to keep up with the literature. Some recent results are very preliminary and require independent validation.

LLM-assisted discoveries: challenges and cryptographic risks in 2026

Stronger classical competition, limits to quantum advantage and threats to cryptography

Raises barriers to quantum advantage | Cryptographic risk | Mixed QC / crypto impact

Date (2026)	LLM / research system	Scientific discovery	Potential impact and caveat
Feb–Mar	GPT-5.2 / internal OpenAI	Single-minus gluon and graviton amplitudes	Analytical simplification of scattering calculations; indirect classical competition.
Jul 28	Claude Mythos Preview	HAWK post-quantum signatures	Improved attack on a PQC candidate; does not generalize to all lattice schemes.
Jul 28	Claude Mythos Preview	Reduced-round AES	200–800× attack improvement on 7 of 10 rounds; full AES not broken.
Sep 24	OpenAI internal model	2D entanglement area law	Supports compact representations of gapped ground states; not efficient contraction.
Sep 24	OpenAI internal model	Polynomial PEPS approximation	Strengthens classical tensor-network baseline for gapped ground states.
Sep 24	OpenAI internal model	Continuum Coulomb QMA-hardness	Ground-state energy can be hard even for quantum computers; state preparation matters.
Sep 25	Claude (version unverified)	Nine-loop scattering amplitude	Extends classical analytical calculations in supersymmetric Yang–Mills theory.
Oct 16	Claude Mythos 5	Reversible elliptic-curve arithmetic	Improved quantum circuits could lower the cost of attacks on ECC.
Oct 26	OpenAI internal model	Exponential Time Hypothesis (ETH)	Claimed subexponential 3-SAT algorithm; could weaken classical complexity baselines.
Nov 06	Anthropic internal model / PSI	Aaronson–Ambainis conjecture	Limits generic quantum query advantages on typical inputs; not a circuit simulation theorem.
Oct–Nov 24	OpenAI internal model / AI-assisted work	Matrix multiplication exponent 2.25	Claimed matching bounds; could improve classical linear algebra and simulation.

Dates: first public disclosures. Discovery titles link to sources. Recent claims may require further independent validation.

The effect can be negative for quantum computing. AI-assisted research may improve classical tensor-network methods, matrix algorithms or the understanding of computational complexity. It can also affect cryptography, including post-quantum cryptographic schemes. A result about a complexity class or a mathematical bound does not automatically translate into a practical algorithm or a successful cryptographic attack. The recent claims displayed here should be interpreted with particular caution.

LLM-assisted discoveries: opportunities and neutral results in 2026

Advances supporting quantum computing, QEC and FTQC, or with indirect impact

Supports quantum computing | Indirect or neutral

Date (2026)	LLM / research system	Scientific discovery	Potential impact and caveat
Jun 01	Claude Opus 4.6, GPT-5.3-Codex, Gemini 3.1 Pro	Bivariate bicycle QEC codes	LLM-guided search found hundreds of candidate codes; hardware validation remains.
Jun 23	GPT-5.4-mini / nano	Lifted-product qLDPC codes	New high-rate codes; potential reduction in physical-qubit overhead.
Jun 29	Claude Fable 5	QAOA ring conjecture	Exact depth-dependent approximation ratio; no industrial advantage established.
Aug 01	OpenAI internal Astra	Closest vector problem (CVP)	Stronger worst-case hardness bounds; no direct attack on standardized PQC.
Aug 01	OpenAI internal Astra	Binary and spherical code bounds	Tighter theoretical limits; not a constructive QEC code improvement.
Aug 01	OpenAI internal Astra	Parallel repetition of entangled games	Stronger soundness bounds for quantum verification and nonlocal games.
Oct 16	Claude Mythos 5	DQI reversible decoder	Reported ~12× gate reduction in a decoding primitive; not end-to-end speedup.
Oct 16	Claude Mythos 5	Quantum chemistry circuit primitives	Exploits molecular structure to reduce reversible computation costs.
Oct 26	OpenAI internal model	4D Yang–Mills existence	Mathematical existence result; mass gap unresolved.
Nov 20	Claude (version unverified)	Riemann hypothesis / GRH	Reported proof; no efficient factoring algorithm follows.
Nov 23–24	Anthropic / OpenAI internal models	VP versus VNP separation	Claimed permanent lower bounds support exact sampling complexity arguments.

Dates: first public disclosures. Discovery titles link to sources. Recent claims may require further independent validation.

AI may also help discover better quantum error-correcting codes, improve reversible circuits and reduce the cost of algorithmic primitives. The net effect on quantum advantage is uncertain because both quantum and classical approaches are improving simultaneously. We cannot assume that the classical state of the art will remain fixed while quantum hardware matures.



https://x.com/plusnineonzero/status/2125653450855969127

The recent announcements and social-media discussions illustrate the speed of this change. There are even rumor alleging that researchers from MIT may have proven that P = BQP. If such a generic equality was proved, the asymptotic case for quantum advantage would be undermined, while practical prefactors might still matter.

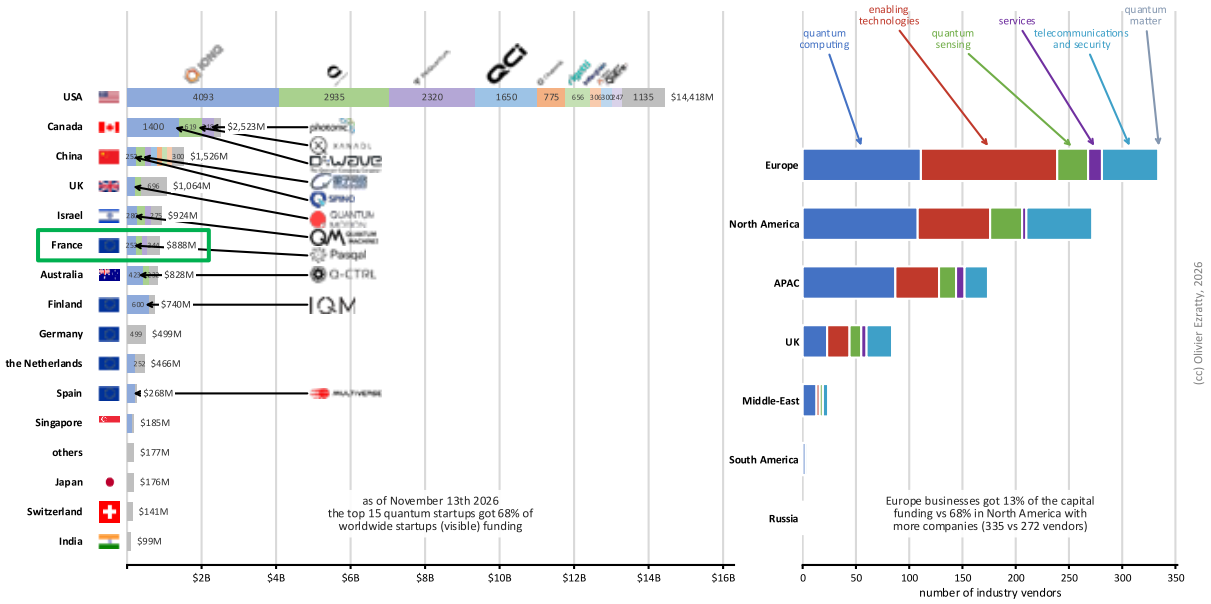

57

the « true » global investments

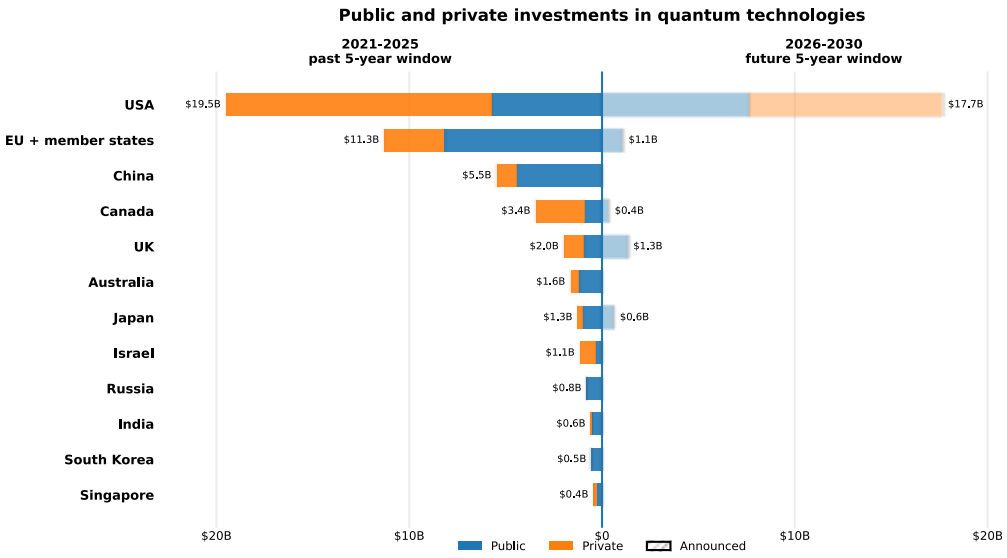

(cc) Olivier Ezratty, 2026

58

The ecosystem is also shaped by the money invested in quantum technologies. Public funding announcements, private capital and actual expenditure are different things. My inventories attempt to distinguish them. The substantial private investment in quantum computing reflects expectations about future markets rather than current commercial returns alone.

industry vendors ecosystem



(cc) Olivier Ezratty, 2026

59

The industrial ecosystem extends far beyond quantum processor manufacturers. It includes software, cybersecurity, sensing, cryogenics, photonics, electronics, manufacturing and materials. These enabling technologies can benefit from the development of quantum systems even before a general-purpose FTQC machine becomes available.

who sells to whom?

China numbers are probably underestimated.

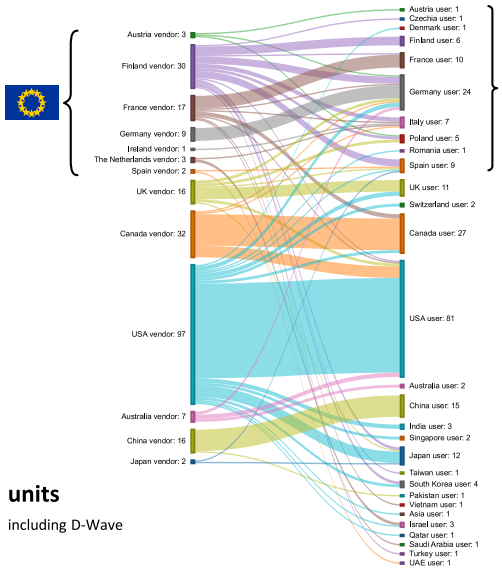

units

including D-Wave

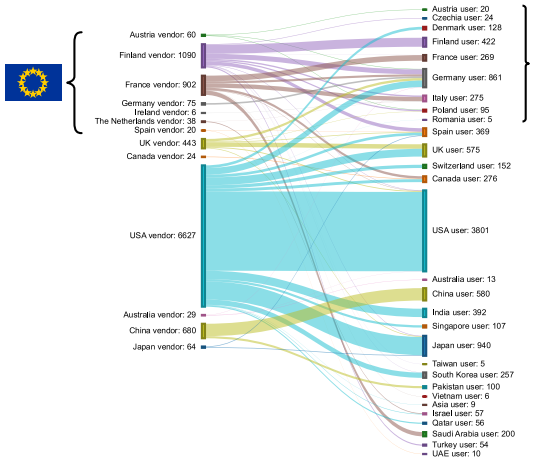

qubits

excluding D-Wave

(cc) Olivier Ezratty, November 2026

I also maintain an inventory of quantum computer shipments worldwide. At the time of this talk, it includes approximately 160 systems. The charts distinguish system units from physical qubit counts. I exclude D-Wave from the qubit-count comparison because annealing qubits are not directly comparable to gate-based physical qubits. The United States leads in physical qubits shipped in this inventory, partly because of IBM. France also has a significant position, partly because Pasqal's analog systems are included. These methodological choices influence the rankings, and Chinese shipment numbers are likely incomplete.

Questions and answers

Question: When quantum computers are integrated with supercomputers, is the energy consumption of the classical infrastructure included?

Usually not in vendor estimates. Classical computing is intrinsic to quantum computing, not simply an optional hybrid configuration. Variational NISQ algorithms use a classical optimizer repeatedly, but even QPE requires classical preparation of the initial state and Hamiltonian encoding. For a large problem, this may require substantial HPC resources.

There is also a practical question of co-location. Quantum hardware can be sensitive to electromagnetic disturbances, vibration and temperature variations, so it may need to be isolated from conventional HPC equipment. A fast classical network between the two may be sufficient for many workloads. The bandwidth and latency requirements depend on the kind of interaction involved, particularly whether real-time feedback is required. Hybrid computing is therefore a combination of hardware, software and connectivity problems.

Question: How do we differentiate quantum and classical solutions, and identify a useful application?

A quantum solution uses a quantum processor as part of the solution, whereas a classical solution uses conventional computing resources. The more interesting question is how to identify a problem for which the quantum component adds value. Sometimes technology developers have a hammer and are looking for a nail, while users start from a problem that they actually need to solve.

I recall discussing this with Bill Phillips, who received the Nobel Prize in Physics in 1997. He was confident that, whatever happened to the commercial market, quantum computers would be useful for solving quantum physics problems. That is a perfectly reasonable scientific perspective, but given the amount of money invested in the industry, we would also like to see a broader societal impact.

End users often create small teams to investigate algorithms and experiment with NISQ hardware. What I find frustrating is when a paper stops at a demonstration on a handful of qubits without estimating the resources required for an application at a meaningful scale. A resource estimate is much more informative because it indicates what the hardware must eventually support.

A serious comparison must also use the best available classical algorithms and hardware. Comparing a costly quantum system with an old classical implementation running on a laptop can create a misleading advantage. I would prefer to compare the quantum solution with an optimized classical implementation on current hardware, and also compare the computing resources available for an equivalent investment.

Question: If the potential impact is uncertain and vendors are not meeting their original goals, why do investors continue investing?

Investors are exposed to different market narratives. My perspective may be more cautious than the prevailing discourse, which often presents quantum computing as an imminent revolution. Companies need to communicate an attractive vision to raise money, and this creates a tension between scientific uncertainty and commercial expectations.

There is nevertheless a positive side. Industrial investment advances science and engineering. A small academic team generally cannot build an entire quantum computing system. Companies such as IBM, Google and specialized startups can undertake full-stack engineering and discover problems that are invisible when each component is studied separately. These problems feed back into research in physics, materials, classical computing, error correction, control electronics and lasers. The development of FTQC is somewhat like a moonshot project. Even if its ultimate commercial impact is uncertain, it can produce valuable scientific and engineering advances along the way.

Question: Where can I find references on FeMoco and quantum chemistry?

My book Understanding Quantum Technologies is freely available on my website, www.oezratty.net, and on arXiv. Its section on quantum chemistry and QPE includes references on FeMoco, including work by Garnet Chan and other researchers. The most recent classical results mentioned in the talk may not yet be included in the edition available at the time of this seminar.

Question: Which quantum hardware modality should a quantum chemist choose?

The first answer is whichever technology becomes available first with the required logical resources. If several platforms can run the relevant circuit, the choice will depend on computing time, cost, accuracy and possibly energy consumption. In a mature FTQC market, I do not expect an absolute separation of hardware modalities according to application type, since many algorithms require similar capabilities: sizable circuits, universal logical operations and adequate error protection.

There can nevertheless be workload-dependent differences. A chemistry calculation requiring many independent circuit repetitions could be distributed across several identical processors. For an algorithm with very few repetitions and a long sequential circuit, the speed of an individual processor becomes more important. Parallelism, connectivity and gate speed thus influence the economics differently depending on the algorithm.

Question: What could quantum computing bring to nonlinear PDEs, particularly Navier–Stokes equations and fluid mechanics?

Nonlinear systems are particularly difficult to handle. Quantum unitary evolution is linear, although measurement and conditioning can be used as parts of algorithms that implement more complex transformations. Approaches include discretization, Carleman linearization and hybrid methods that separate linear and nonlinear terms, solving some parts classically and others on a quantum processor.

The important question is where the end-to-end benefit comes from. There is a significant body of theoretical literature, but many papers stop at the mathematical formulation without providing an implementation in gates or a complete resource estimate. For realistic fluid mechanics, including fire propagation or aeronautical applications, data loading, nonlinear dynamics, precision and the amount of information to retrieve can all be major challenges. I am not sufficiently specialized in these numerical methods to identify a definitive solution. It remains an active area of investigation.

Question: Are industrial users hiring specialists to develop quantum applications for their own businesses?

There are several categories of players. Large end users such as financial institutions and energy companies have established internal teams, sometimes including doctoral researchers, to investigate their own use cases. Specialized quantum software companies develop algorithms and sometimes resource estimates for their customers. Consulting and IT service companies have created teams combining their existing industry expertise with quantum computing. Established software vendors such as Dassault Systèmes are also exploring the field.

Quantum software startups frequently recruit physicists or quantum chemists because they are familiar with Hamiltonians, quantum dynamics and the mathematical formulation of the problems. These skills complement conventional computer science and software engineering skills.

Question: Could better materials reduce quantum error correction overhead and energy consumption?

Yes. In superconducting qubits, researchers are investigating alternative materials for resonators and Josephson junctions, including tantalum and different barrier materials. Material choice, elemental purity, interfaces and deposition techniques can all influence noise. Isotopic purity is particularly relevant to silicon spin qubits, where silicon-28 reduces the effect of nuclear spins.

The benefits of a particular manufacturing process are not always intuitive. Epitaxial deposition does not automatically yield better qubit performance than amorphous deposition. One needs to understand the defects and noise mechanisms associated with the actual materials and interfaces. Research at institutions such as imec illustrates the importance of experimentally comparing fabrication techniques.

The connection with energy is indirect but important: better physical fidelities can reduce the resources required for error correction and therefore the power and energy needed for a computation. Materials science also matters in enabling technologies such as lasers. The stability and longevity of laser materials can affect the reliability of an atomic quantum computer during a long computation.