



Can we mitigate the quantum hype side effects?

We are in the midst of quantum hype with some excessive claims of quantum computing potential, many vendors' and even some research organizations' exaggerations, and a funding frenzy for very low technology readiness level startups. Governments are contributing to this hype with their large quantum initiatives and their technology sovereignty aspirations.

Technology hypes are not bad per se since they create emulation, drive innovations and contribute to attracting new talents. It works as scientists and vendors deliver progress and innovation on a continuous basis after a so-called peak of expectations. It fails with exaggerated overpromises and underdeliveries that last too long. It can cut short research and innovation funding in the mid to long term creating some sort of quantum winter.

In the "Mitigating the quantum hype" essay published on Arxiv on February 7th, 2022 (26 pages), I try to paint a comprehensive picture of this hype and its potential outcomes. This piece is a abridged version of the paper.

It can also happen when some technology innovation is not creating real perceived societal progress and value as has, when technology study success, its visibility shows back with a growing "slope of enlightenment" and a "plateau of productivity". The technology has then the potential to become mainstream and commoditized in the marketplace, whether in the enterprise or the consumer space depending on the technology.

The Gartner hype curve is an over-simplified and empirical model that does not capture well what can happen with disruptive science and technologies". It is constructed with adding a natural "hype" half curve and a technology maturity S curve, given these do not rely on the same metrics. It is missing scientific, technology and even business rationale. It deals mainly with the emotional response to new science and technology and is not of any assistance. It also presupposes that a hype's visibility, expectations and needs are correlated. Its premise is that at any point in time, the given technology domain will succeed. But what if science and technology fail to address? What if the new needs the innovation is supposed to address do not really exist? The curve is also built out of a strong negative bias, fearing the various technologies that really failed for one reason or the other".

The market is now flooded with a growing number of technology trends and fads. How is this affecting visibility? Sometimes, the technology trend is not well defined, the "hypeability" which drives the market "copy". The length of the trough of disillusionment can be very long spanning several decades. We are still in it for mobile AI and quantum several decades. What if the new needs the innovation is supposed to address do not really exist? The curve is also built out of a strong negative bias, fearing the various technologies that really failed for one reason or the other".




Figure 1 - Gartner hype curve model and some alternate realities.




Figure 2 - Geoffrey Moore's innovation adoption model.

If we were to use Moore's model with quantum computing, we'd rather say that we are still in the innovation stage with existing pre-NSQ quantum processing systems that are still experimental devices". The business feature set entrepreneurs to reach some quantum advantage to solve practical business problems. Early adopters will jump on the bandwagon with useful NSQ or large scale and full-fledged computing systems. It is underdeveloped hardware and then, when the software and cloud ecosystem grows, we'll cross the chasm and see the early majority adopt quantum computing massively.

Christman's innovation disruption model covers innovations that are spreading to lower or underserved consumers and then become mainstream and consequently are based on innovative business models". The model works when a new technology and its implementation significantly expand a given market. The classical examples are the personal computer and the smartphone. It is less applicable to quantum computing given it will, at best at the beginning, be a sub-market of the mature high-performance computing market.

In both cases, Christman's and Moore's model are describing the implementation details of the slope of enlightenment in the Gartner hype model, when technology starts to actually work. In the case of quantum computing, these considerations will become useful when quantum computers really scale and provide some quantum advantage compared to classical computers. But some aspects of Moore's innovation adoption model are just in place only on, such as software tools, ecosystems and platform like IBM Qiskit. It prepares the market, only on, for a dissemination of software solutions and skills when scalable quantum hardware actually starts to work, and can become a leading platform.

Symbolic artificial intelligence

In the 1980s, symbolic AI and expert systems were trendy but had practical implementation issues. Capturing expert knowledge was difficult and could not be automated which caused very an uncertainty regarding its future are today. There were even dedicated machines tailored for training the artificial intelligence LISP programming language.

I found that, although there is some significant uncertainty on the potential to create real scalable quantum computers, the scientific and vendor fields are relatively sane and solid compared to other technology hypes. What is striking is how the vendors hype has some profound and disruptive impact on the organization of fundamental research.

I then make some proposals to mitigate the potential negative effects of the current quantum hype including recommendations on scientific communication to strengthen the trust in quantum science, vendor behavior improvements, benchmarking methodologies, public education and putting in place a responsible research and innovation approach.

History lessons and analogies

Artificial intelligence specialists who have been through its last “winter” in the late 1980s and early 1990s keep saying that quantum computing, if not quantum technologies on a broader scale, are bound for the same fate: a drastic cut in public research spendings and innovation funding. Their assumption is based on observing quantum technology vendors and even researchers overhype, on a series of oversold and unkept promises in quantum computing and on the perceived slow improvement pace of the domain.

The industry payback of quantum physics research came with the transistor invention in 1947, the laser in the early 1960s, and many other technology feats (TVs, LEDs, GPS, ...) leading to the “first quantum revolution” based digital era we are enjoying today. The “second quantum revolution” era deals with controlling individual quantum objects (atoms, electrons, photons) and using superposition and entanglement. The early 1980s were defining moments with Yuri Manin and Richard Feynman expressing their ideas to create respectively gate-based quantum computers and quantum simulators (in 1980 and 1981). We are now 40 years ahead, and even though quantum science advances have been continuous, usable quantum computers offering a quantum advantage compared with classical computers are not there yet and this can be the source of some impatience. However, three other applications of this second quantum revolution are alive and well: quantum telecommunications, cryptography and sensing. To some extent, the latter is well under-hyped.

The last decades saw an explosion of digital and other technology waves, most of them successfully deployed at large scale. Micro-computers invaded the geek world, then the workspace and at last, our homes. The early 2000s marked the beginning of the consumer digital era, starting with web-based digital music, digital photography, digital video and television, e-commerce, the mobile Internet, the sharing economy, all sorts of disintermediation services and at last cloud computing. There were however some failures, with hype waves and adversarial outcomes. In many cases, while these hypes peaked, some skepticism could be built out of common sense.

In the 1980s, symbolic AI and expert systems were trendy but had practical implementation issues. Capturing expert knowledge was difficult and could not be automatized when contents were not massively digitized as they are today. There were even dedicated machines tailored for running the artificial intelligence LISP programming language. It could not compete with generic CPU based machines from Intel and the likes that were benefiting from a sustained (Gordon) Moore’s law. As a result, the couple related startups failed even before they could start selling their hardware. Then, expert systems went out of fashion. AI rebirth came with deep learning, an extension of machine learning created by Geoff Hinton in 2006. The comeback can be traced to 2012 when deep learning could benefit from improved algorithms and powerful GPUs from Nvidia. The current AI era is mostly a connectionist one, based on machine and deep learning probabilistic models, as opposed to symbolic AI based on rules engines. As an extension of the old symbolic AI winter, autonomous cars are also progressing quite slowly. The scientific and technology challenges are daunting, particularly when dealing with the interaction between autonomous cars and humans, whether they are walking or in other 2, 3, 4 or more wheels vehicles.

Nowadays, artificial intelligence is an interesting field with regards to how ethics, responsible innovation and environmental issues are considered, mostly as an afterthought. The learning here is, the sooner these matters are investigated and addressed, the better.

The famous Gartner hype curve can move ahead or stop after the trough of disillusionment. When it stops, it can be related to some societal and economical reason (like there is no real business case, people do not like it, it is not a priority, it is too expensive) or to scientific showstoppers (it does not work yet, it is unsafe, etc.). In the science realm, where are the overpromises and hypes that led to some sort of technology winter?

You can count nuclear fusion in, and it is still a long-term promise, although being back in fashion with a couple startups launched in that domain. These have raised more than \$2B so far. Being able to predict when, where and how nuclear fusion could be controlled to produce electricity is as great a challenge as forecasting when some scalable quantum computer could be able to break 2048 bit RSA keys. In a sense, these are similar classes of challenges. Unconventional computing brought its own fair share of failed hypes with optical computing and superconducting computing. We have been hearing for a while about spintronics and memristors that are still in the making. These are mostly being researched by large companies like HPE, IBM, Thales and the likes.

These cases are closer to the quantum hype in several respects: their scientific dimension, fact-checking toughness, low technology readiness levels and scientific uncertainty. They are also driving a lower awareness in the general public than non-scientific hypes like cryptocurrencies, NFTs and metaverses.

Quantum hype characterization

The current quantum computing hype started in 1994 when Peter Shor, then at the Bell Labs and now at the MIT, created his famous eponymous integer factoring quantum algorithm that could potentially break most public-key-based encryption systems. Shor's algorithm created some excitement with governments, particularly in the USA. The potential to be the first country to break Internet security codes was appealing, as a revival of the famous Enigma machine during World-War II. At least, when not taking care of its potential impact on society as a whole.

As the web and mobile Internet usage exploded, starting in the mid to late 2000s, the interest in quantum technologies kept growing. The Snowden 2013 revelations on the NSA Internet global spying capabilities had a profound impact on the Chinese government. It triggered or amplified their willingness to better protect their sensitive government information networks. It may have contributed to their deployment of a massive quantum key distribution networks of fiber-optics landlines and satellites.

In the early 2000s, research labs started to assemble a couple of functional qubits, starting with 2 to 5 qubits between 2002 and 2015. We now have up to 127 superconducting qubits in the USA (IBM), but these are way too noisy to be practically usable together. Many challenges remain to be addressed to create usable and scalable quantum computers, one of which involving assembling a large number of physical qubits to create logical "corrected" qubits. This engenders a large overhead creating significant scalability challenges both at the quantum level and at the classical level, with electronics, cabling, cryogenics and energetics. It is a mix of scientific, technological and engineering challenges, the scientific part being stronger than with most other technologies. Other researchers and vendors are trying to exploit existing noisy systems (NISQ, noisy intermediate scale quantum) with noise-resistant algorithms and quantum error mitigation. There is not yet a shared consensus on the way forward. To some extent, it is good news. We still need some fault-tolerance in the related scientific and innovation process.

Still, nowadays, the quantum computing frenzy is unbridled. The perceived cryptography threat from quantum computing and Shor's algorithm is still at its height. It drove the creation of post-quantum cryptography systems (PQC) running on classical hardware. NIST is in the final stages of its PQC standardization process which should end by 2024. Many organizations will then be mandated to massively deploy PQCs. We are in a situation equivalent to those folks who built their own home nuclear shelters, reminiscent of some advanced form of technology survivalism.

Some classical aspects of hype can be found with quantum computing.

First, quantum computing use cases promoted by vendors and analysts are far-fetched. The classics are ab-

initio/de-novo cancer curing drugs and vaccines development, solving earth warming with CO₂ capture, creating higher-density, cleaner and more efficient batteries and Haber-Bosch process improvements to create ammonium more energy-efficiently. Of course, all these Earth-saving miracles will be achieved in a snap, compared to billion years of classical computation. Most of these promises are very long term, given the number of reliable (logical) qubits they would require.

Second, the quantum computing scene is prone to fairly inconsistent signals, particularly for appreciating the real state of the art. On one hand, research labs and vendors tout so-called quantum supremacy and advantages, showcasing quantum computers that are supposedly faster than the best supercomputers in the world. On the other hand, we also hear that “real” quantum computers do not exist yet and could appear only in about 10 to 20 years. Reconciling these contradictory messages and fact-checking these claims is far from being easy. As a result of these mixed and conflicting signals, many non-specialists extrapolate these fuzzy data and make rather inaccurate, particularly overoptimistic, interpretations of quantum computers readiness for prime-time usage, like when a quantum computer will be able to break an RSA-2048 encryption key and threaten Internet’s whole security.

Third, like all digital technology gold rushes, market predictions are overly optimistic. The quantum computing market is supposed to reach \$2,64B in 2022 for Market Research Future (in 2018). Then we have \$8,45B in 2024 for Homeland Security (in 2018), \$10B in 2028 for Morgan Stanley (as of 2017), \$15B by 2028 for ABI Research (2018) and \$64B by 2030 for P&S Intelligence (in 2020). As of early 2020, McKinsey predicted that quantum computing would be worth one trillion dollars by 2035 (i.e. \$1000B). This forecast’ bias comes from a trick used a few years ago to evaluate the size of Internet of things, artificial intelligence and the hydrogen fuel markets. It is not a market estimation for quantum technologies as such, but the incremental revenue and business value it could generate for businesses, such as in healthcare, finance or transportation verticals. The \$1T McKinsey market sizing has been exploited again and again in the media and even by financial scammers. One can wonder whether McKinsey anticipated that side effect.

Fourth is what’s happening in the vendors scene with some startups racking-in hundreds of millions of dollars of funding. This funding trend can be explained by several phenomena: the lack of funding of ambitious consolidated research programs in the public space, the willingness of quantum researchers to create their own venture with more freedom, the abundance of money in the entrepreneurial scene and the quantum hype itself. It creates a very unique new situation. Due to their low technology readiness level (TRL), quantum computing startups are mostly private fundamental research labs who also happen to undertake some technology developments. In the classical digital world, most startups funding cover product developments, industrialization, ecosystems creation and, above all, customer acquisition and housekeeping costs. The resulting scales are mind-blowing for academic researchers. Startups raising about \$20M can create research teams that are bigger than most publicly funded research teams with over 50 PhDs, and with real well paid jobs, not short-term post-docs tenures. The couple startups who raised more than \$100M created teams with hundreds of highly skilled researchers and engineers. You also have to boil in the significant quantum R&D investments from the large US IT vendors which I’ll nickname the “IGAMI” (IBM, Google, Amazon, Microsoft, Intel). They play a similar role to the Bell labs in the telecom industry, a couple decades ago.

In some regions, this phenomenon can create a huge brain drain on skilled quantum scientists. It drives structural changes on how quantum computing research is organized. Academic research organizations now have a harder time attracting and keeping new talent. In many countries where there are not many opened full-time researchers jobs positions, PhDs and post-docs who are looking for a well-paid and full-time job are obviously attracted by the booming private sector.

As a consequence, many scientists feel the pain. They observe, mostly silently, all the above-mentioned claims and events and can feel dispossessed. They know scalable quantum computing is way more difficult to achieve

than what vendors are promising. They fear that this overhype will generate some quantum winter like the one that affected artificial intelligence in the early 1990s after the collapse of the expert systems wave. Scientists made entrepreneurs feel obliged to leave good scientific methods and practices aside and, for example, favor secret against openness, and as, a result, avoid peer reviewing. This is one of the reasons why it is so difficult to objectively compare the various early quantum computers created by industry vendors, from the large ones to startups with a common set of benchmark practices.

Quantum hype specifics

How different is the quantum hype phenomenon compared to the typical other hypes There are many differences, mostly related to the scientific dimension and diversity in the second quantum revolution.

The largest difference is the nature of the uncertainty that is mainly a scientific one, not a matter of market adoption dynamics or some physiological limitations. It is even surprising to see such hype being built so early with regards to the technology maturity, at least for quantum computing. This is an upside-down situation compared to the crypto world: there is market demand but a low technology readiness level (in quantum computing) versus a moderate market demand and a relatively high TRL (in cryptos).

The second is that quantum physics is serious research and less prone to fakery than life science and human sciences. The field has only a few retracted papers in peer reviewed scientific publications. A survey published by Nature in 2016 did show that chemistry, physics and engineering were the fields where reproducibility was the best while life science is the worst. There are however some cases when quantum research findings cannot easily be reproduced because they are very theoretical. It happens with many quantum algorithms which can only be either emulated classically at a very small scale or only performed at an even smaller scale on existing qubits. The same can be said of new quantum error correction codes. And on top of that, getting some exponential acceleration with quantum computing requires the creation of large entangled sets of qubits, which has not been achieved yet beyond 30 qubits.

The third difference is that quantum technologies are extremely diverse. There are many competing technology routes, particularly with the types of qubits that could be used to build scalable quantum computers. Although it could be perceived as slow, scientific and technology progress in quantum computing is steady. On top of that, there are several competing quantum computing paradigms: gate-based quantum computing, quantum annealing and quantum simulation. This creates a sort of fault-tolerance in the quantum computing ecosystem. Quantum simulation systems like the ones from Pasqal may bring some computing advantage earlier than imagined, way before gate-based systems. Also, although not a darling in the quantum computing world, quantum annealing and D-Wave are making regular progress and are publicizing the greatest number of potential use cases for their systems.

However, the scientific and technology understanding levels required to forge some view on the credibility of scalable quantum computing are very high. It is the same in other fields like quantum telecommunications and cryptography. It is higher than any other discipline I have investigated so far, out of the above-mentioned ones in other technology fields.

The current quantum hype is also fueled by some regions and countries who categorized quantum technologies in the field of sovereign technologies. These are technologies a modern state must either be able to create or to procure. Like with many other strategic technologies, geographical competition happens at least first between the USA and China, with Europe then sitting in between and willing to get its place under the quantum sun. Quantum technologies are considered as being “sovereign” for a couple reasons. The first relates to cybersecurity. Shor’s algorithm potentially threatens the cybersecurity of most of the open Internet. The second is the European Union is traumatized by its dependence on the USA and Asia for many critical digital

technologies. Europe is eager to get its share with quantum technologies, leveraging a potential opportunity to level the playing field with a new technology wave, having missed most previous digital waves. The third is “dual usages”, where quantum technologies can be applied in the civil and commercial spaces and in the military/intelligence. The fourth relates to securing access to special raw materials and isotopes as well as to some critical enabling technologies.

These national security concerns tend to exaggerate threats, thus fueling the hype. It is prone to propaganda, the strongest coming from China and its various technology premieres. It also tends to limit international cooperation and as a result reduces the capacity to fulfill expectations that would require a coordinated worldwide mobilization. Thus, there is a need to differentiate what belongs to worldwide shareable fundamental research and where it becomes strategic and not shareable, especially given these concerns are also emerging as quantum technologies enter the commercial vendors space.

Quantum hype mitigation approaches

Technology hypes are inevitable. But there are ways to mitigate their most undesirable effects like a sluggishness to deliver results or a decrease of trust in scientists, engendering some winter that would shrink or kill government, investor and corporate R&D funding sources. It is a matter of general conduct for the whole quantum technologies community but also some specific actions.

A first ongoing action is to educate the public and explain the progress. Although scalable quantum computers have not yet seen the light of day, quantum science is making a lot of progress. Qubit fidelities are improving in many of their implementations. What is amazing in the quantum computing space is the sheer diversity of technologies being developed.

A second would be to better map the challenges ahead. Quantum computing critics must be listened to. It is a mix of theoretical, scientific and engineering challenges, the best example being the links between quantum thermodynamics and cryo-electronics. But don't count on the magic wand of engineering to solve quantum computing' daunting scientific challenges.

A third one would be to consolidate some key transversal projects. Due to its diversity, quantum technology efforts are far from being concentrated or consolidated, even in the USA and China. The efforts are scattered in many laboratories and vendors, since many different technology avenues are being investigated. Many physicists argue that it is way too early to launch consolidated initiatives given the lack of maturity of qubits physics. There are some domains where it would make sense like with cryo-electronics, with scale-out architectures connecting several quantum processing units together through photonic links and with quantum technology energetics. Can we ensure quantum computers will bring some energetic advantage and not see their scalability be blocked by energetic constraints? These are perfect cases where multidisciplinary science and engineering are required. Hardware and software teams would also benefit from working more closely together. At some point, creating a fault-tolerant quantum computer will look more like building a passenger aircraft, a nuclear fusion reactor ala ITER instead of a creating a small-scale physics experiment. There, much more international coordination and collaboration will be required to assemble all of it in a consistent manner.

A fourth approach would be about creating rational ways of comparing quantum computers performance from one vendor to another and across time is fundamental to set the stage. IBM's Quantum Volume was a first attempt, launched in 2017 and adopted in 2020 by Honeywell and IonQ. It is however a bit confusing, particularly given quantum volume cannot be computed in a quantum supremacy or advantage regimes since it requires doing some emulation on classical hardware. Having vendors accepting the implementation of third-party benchmark tools is important to create some form of trust. These benchmarks will lead to the creation of quantum computers rankings like the Top500 HPC. At some point, we will also need to create energetics

oriented benchmarks and rankings, to ensure quantum computers scaling is not done at the cost of astronomical energy costs.

A last, mitigating the quantum hype is about the addressing science and technology ethics requirements. It should not be an afterthought, studied only when technology is ready to flood the market. It starts with studying the potential ethical and societal issues created by quantum technologies with Technology Assessment methodologies and can then lead to Responsible Research and Innovation approaches that take into account ethics, societal values and imperatives early on, like fighting climate change. It tries to anticipate the unintended consequences of innovation before it occurs, even though innovation processes are inherently messy and unpredictable and when the perspectives of quantum technologies are polluted by off-the-mark wild claims. It can also lead to create new regulations such as on privacy. It should also include the matter of resources like raw materials and energy. These are unavoidable constraints, particularly as materials are finite and electricity production has still a significant CO₂ emission cost.

Useful quantum computing is probably bound to be at best a long lasting quest, spanning several additional decades. Patience will be a virtue for all stakeholders, particularly governments and investors.

Cet article a été publié le 9 février 2022 et édité en PDF le 9 février 2022.
(cc) Olivier Ezratty – “Opinions Libres” – <https://www.oezratty.net>