



Universal
Quantum



Scalability is the lifeline of quantum computing

UQ Insight Report

Table of contents

Executive summary	3	Closing thoughts	13
Unlocking the potential of quantum computing	5	A thought for investors	13
Humanity needs greater computational capacity	5	A thought for governments	13
Developing computing solutions that complement conventional computing	5	References	15
Quantum computing will be a gamechanger	5		
Quantum computing invites us to dream big one more time	5		
Quantum computing has not delivered broad utility yet	5		
Noise is killing quantum computers	6		
Current best quantum computers fall short by eight orders of magnitude	6		
Error correction is a must	6		
We must develop quantum computers with millions of high-quality qubits to succeed	7		
It is not just about current qubit count and current error rates	8		
Performance is more than just gate speed	8		
Scalability is the lifeline of quantum computing	9		



Executive summary

Threats to humanity

Ten million people die from cancer each year, and over 50 million people suffer from Alzheimer's. Climate change marches on relentlessly, inequities around the world are ever-increasing, and threats to our global peace and security are accelerating. It is now imperative that we increase our computational capacity to address these challenges in order to preserve and elevate all living systems. Quantum computing has the potential to play a key role in this. It is a new compute technology that tames the strange effects that come to the fore on the microscopic scale to provide new capabilities for solving some of humanity's toughest computational challenges.

The truth of where we are

Despite optimistic aspirations and statements over many years, let's be clear: quantum computing has not delivered broad utility that helps alleviate any of humanity's challenges—yet. Facing the truth about this reality, as well as embracing the path to reach the ultimate goal, is critical to hold in focus for governments and the commercial sector to make the decisions that will ultimately ensure we deliver for humanity.

What is holding back quantum computing: noise

In a nutshell, disturbances to the accuracy of calculations, or "noise," are leading to errors in quantum computers that limit their usefulness. These currently available machines are often referred to as Noisy Intermediate-Scale Quantum (NISQ) computers. The current best-

performing NISQ machines, operating with tens of qubits, fall short by up to eight orders of magnitude from the error rates required to solve some of the most impactful applications [1, 2, 3, 4].

The solution i: error correction

The way out of this is error correction. Error correction uses multiple high-quality qubits to deliver a logical qubit. When done right, this logical qubit can have a significantly lower error rate than its underlying physical qubits. Generally, the more qubits and the lower the error of the underlying qubits, the lower the error of the logical qubit. Many logical qubits are then needed to make a useful quantum computer.

The solution ii: a fault-tolerant quantum computer (ftqc)

When mapping this to the high-impact applications of quantum computing, we find that in practice, the goal is an FTQC with millions of high-quality physical qubits. This is the ask, the goal, the absolute requirement to enable quantum computing to deliver for humanity.

Identifying the best approach

So how do we identify the best approach towards reaching that goal? It turns out that one of the most popular "tests"—asking, "What is the current number of qubits, and error rate, of the quantum computer?"—gives only a partial picture of the architecture. It tells us little about factors such as modularity, cost, manufacturability, repeatability, and controllability, all of which are essential ingredients to achieve scalability. Scalability

is the key factor to achieve useful quantum computing, so we must ask questions that uncover the promising paths towards achieving that goal.

The role of venture capital and the commercial sector: short vs. long term

The capital markets play a crucial role in developing new technologies. Venture investment is playing an active role in quantum computing; however, many of these investments have limited long-term gain given the inherent limitations of NISQ. These commercial driving forces do not necessarily work for FTQC, which seeks sustained, long-term patient capital investment required to deliver the true promise of quantum computing. The widespread short-term thinking of both the commercial and venture sectors is driving quantum computing development to focus on short-term gains. There is an ever-decreasing overlap between those short-term gains and the developments needed to build an FTQC—a risk not only for the technology itself but also the investments.

The critical position for government

The inability of established commercial driving forces to push forward the development of an FTQC as quickly as it could is a huge risk governments must consider. Loss of a technical sector, price points for access to non-domiciled technology, capital flight to successful regions, industrial and academic decline, and supply chain disablement. Ceding quantum capability could have a detrimental impact on a nation's industry and economy. We must also be mindful that certain state actors are likely seeing an opportunity to leapfrog into

a leadership position (in both defence and commercial sectors) by focusing efforts on developing an FTQC as quickly as possible.

Our insight report

It is not possible here to provide a complete picture of quantum computing and include all the nuances, but if you want to dive a bit deeper, then our report is a good place to start. It includes:

- More details on the journey ahead for quantum computing
- Important questions to ask when evaluating the viability of particular architectures for reaching the all-important FTQC era
- Some thoughts for investors and governments

Unlocking the potential of quantum computing

Humanity needs greater computational capacity

Computing has transformed our lives in ways unimaginable; nonetheless, humanity's progress and well-being are hampered by the limitations of our current computing capabilities that are not immediately obvious, especially in high-impact areas that include medical research, climate change, AI, global peace and security, and material science. We live in a world where 10 million people die from cancer each year and over 50 million people suffer from Alzheimer's, where climate change marches on relentlessly, and where threats to global security are accelerating. We must increase our capacity to find answers to these challenges.

Developing computing solutions that complement conventional computing

Developing computing solutions that complement conventional computing is, therefore, a critical challenge in moving humanity forward and, crucially, in protecting and bettering people's lives and their environment.

Quantum computing will be a gamechanger

Quantum computing is one of the most exciting computing technologies currently being developed to help address these challenges. Remarkable strides have been made since the early academic work demonstrating exquisite control over quantum systems such as atoms and photons, and the ability to turn them into functioning qubits — analogous to the “bits” in

conventional computing. Relentless curiosity and ingenuity have since transformed this early work into prototype quantum computers that can conduct key operations fundamental to quantum computing, and these machines nowadays use qubits made up of a range of different quantum systems, all at various levels of technological readiness level (TRL). Crucially, our understanding of the applications of quantum computing is progressing immensely and will help to formulate solutions to the challenges humanity is facing.

Quantum computing invites us to dream big one more time

We cannot yet precisely say to what extent quantum computing will be able to answer the many challenges facing humankind, as well as the ecosystem we live in, but we do already know it will play a huge part. Just imagine a world where we have more affordable drugs for some of the worst diseases, a world where we have answers to some of the greatest scientific mysteries, a world where humanity has access to a tool of unimaginable computational power at their hands that can be used to make their biggest dreams a reality, a more equitable, happy, and healthy world. Quantum computing will be a key driver of making that world a reality. It is worth remembering that we did not have any appreciation of the full impact conventional computing would have at the beginning of that journey, and look where computing has taken us! Quantum computing invites us to dream big one more time.

Quantum computing has not delivered broad utility yet

The excitement around quantum computing has led to a flourishing and vibrant commercial

quantum sector where access to prototype quantum computers is now only a click away, and announcements on the next 'big breakthrough' are almost ubiquitous.

Despite all the noise, hype, and a huge amount of effort, quantum computers available today, which are generally referred to as Noisy Intermediate-Scale Quantum (NISQ) computers, have not yet reached the point of being able to deliver broad utility that addresses some of the big challenges of humanity; and in reality, NISQ computers most likely never will.

Noise is killing quantum computers

So, what is the problem? All quantum computers suffer from noise that destroys the qubits, noise from the environment, and noise that we introduce when controlling the qubits. This noise leads to errors that dramatically limit the time and number of operations you can run on a quantum computer and fundamentally limits its usefulness. Because of these errors, adding more qubits does not improve the performance of a NISQ computer unless you also reduce the errors.

Current best quantum computers fall short by eight orders of magnitude

Some of the most high-impact applications require error rates of around one error in every trillion operations (that is not a typo) in a quantum computer that has a few thousand qubits [1, 2, 3, 4]. Current best error rates for the most important quantum operations currently sit between one error in every thousand and ten thousand operations on tens of qubits. We are short by over eight orders of magnitude, and considering it took the field decades to get where we are right now, it is a somewhat

insurmountable challenge to get to the error rate required without actively addressing the errors.

Efforts are currently ongoing where the big development goal centers around improving the currently best error rates by an order of magnitude or so and adding a few more qubits to try and unlock a low-hanging application that may add some limited value, but this is far removed from creating the sort of utility humanity ultimately demands.

Error correction is a must

So, what is the solution? Given that all qubit implementations, including the high-quality ones, have noise that is orders of magnitude too high to run the most impactful algorithms, we must actively correct these errors programmatically. There are multiple avenues being pursued; here is a high-level summary of approaches to dealing with errors:

- Algorithms that use short reconfigurable (or variational) circuits on current small-sized quantum computers use classical computing to lower the effective error rate [5]. These "NISQ" techniques are not scalable as the size of the algorithm (or circuit) that can be executed is limited.
- Running repeated instances (shots) of the same algorithm; this is known as Quantum Error Mitigation. The number of shots rises exponentially as the size of the algorithm gets larger, again making it unsuitable to work at scale.
- Quantum Error Correction (QEC) techniques use many physical qubits to create a "logical" qubit that can have vastly superior error rates compared to the many "physical"

qubits that are used to create them. Generally speaking, the more physical qubits used per logical qubit, the lower the error on the logical qubit. By using hundreds or more physical qubits per logical qubit, we can reach sufficiently low error rates to run the most powerful applications.

QEC is, therefore, the method that ultimately gets us to the low error rates we must get to.

It is important to note that simply creating a logical qubit does not mean it has a sufficiently low error rate. The quality of the logical qubit is a function of three factors: the underlying physical qubit error rate, the number of physical qubits used to create one logical qubit, and the chosen quantum error correction technique (QEC code). For example, hardware that has access to high-fidelity long-range connections between physical qubits can utilize more efficient QEC codes. Logical qubits demonstrated in the near term will not yet have sufficiently low error rates to run the most impactful quantum algorithms.

We must develop quantum computers with millions of high-quality qubits to succeed

Within this very noisy and confusing environment that is the quantum computing sector, it is worth zooming out to ask how we go as quickly as possible from the incredible work that has been done to date to the goal where humanity uses quantum computing to improve the state of things.

Before diving deeper into answering this question, just bear the following utility metrics in mind:

- Want to simulate large molecules that could, for example, be catalysts for drastically reducing emissions? You need millions of qubits [1, 6].
- Want to accelerate drug discovery by simulating complex protein interactions? You need millions of qubits [7].
- Want to design new materials with enhanced properties, such as superconductors? You need millions of qubits [8].
- Want to break RSA encryption? You need millions of qubits [2].
- Want to perform derivative pricing for the financial sector? You need millions of qubits [9].

We must be clear that quantum computers with millions of individually controllable, error-corrected, and high-quality qubits are required. Such machines are referred to as fault-tolerant quantum computers (FTQC). This need has been well known for a long time; however, it has not been given the necessary focus of many decision-makers when big expectations were raised by many in the field for what could be done with NISQ computers. Valuable resources are being diverted to NISQ computers when they are critically needed for FTQC development.

Developments on scaling to FTQC are currently unfolding in two ways:

- Producing a large number of qubits and then solving the problem of controlling them well enough to implement error correction and then drive the error down low enough to unlock applications.

- Relentlessly lowering the underlying error rate and then working out how to scale this up.

Both approaches have their advantages and disadvantages, and most likely, the winning approach lies in the blend of these approaches. Purely focusing on reducing the error rates and adding a few more qubits while overlooking the important work of actually making sure that the solutions used to enable those lower error rates can scale to millions of qubits adds a huge risk. The same holds true for approaches focusing on just adding a large number of qubits without having good answers for how to get the errors sufficiently low.

It is not just about current qubit count and current error rates

These two metrics say very little about the chance of reaching the million-qubit scale. Having a frank conversation about what quantum computers will be able to do when, and basing this on our scientific understanding, is absolutely crucial. As is a conversation around what is required to scale these machines to millions of qubits, as well as how to evaluate which architecture(s) have a good chance of delivering that. The latter is particularly difficult to assess, as the commonly used metrics to evaluate current quantum computers, such as the number of qubits and fidelities, tell us very little about the scalability of the approach used for achieving those metrics in the near term. For example, a quantum computer with extremely high fidelity today does not indicate its ability to scale to millions of qubits, nor does it provide insight into maintaining those fidelities at that scale. Similarly, the current qubit count tells you

little about the architecture's ability to support orders of magnitude more qubits. This is not to say that these metrics are not important, but additional key questions must be considered when evaluating an architecture's potential to reach the goal of millions of qubits, a milestone crucial for the advancement of humanity and the world. We must delve into those key questions to gain a better understanding of where the future utility of quantum computing really lies.

Performance is more than just gate speed

A quantum algorithm will execute quicker on slower gates with certain characteristics compared to faster gates without those characteristics. Gate speed, i.e., the speed at which quantum gates (analogous to logic gates in conventional computing) can be performed, is commonly used as a metric to indicate how fast an application might perform on a given architecture. Similar to conventional computing, where the CPU clock speed does not ultimately fully inform us about the performance of an application (there are factors like memory bandwidth, network bandwidth, disk read/write speeds, etc., that affect performance), there are many other aspects that affect the performance of a quantum computer.

It is, therefore, worth having a closer look at the key drivers of the performance of a quantum computer, which can be summarized as follows:

- Qubit count
- Fidelity of operations (e.g., quantum gates, input, and output of information, etc.)

- Code cycle time (the code cycle time is made of all time-consuming operations such as those required to control the qubits, move the quantum information around the quantum computer, and input into and extraction from the qubits the required classical information)
- Qubit connectivity (ability to only carry out operations between qubits that are directly next to each other (nearest neighbour connectivity) vs. ability to directly connect qubits 'further away' (all-to-all connectivity))
- Algorithmic efficiency (improving the mapping of problems to quantum algorithms for reduced qubit and operation requirements, and further optimizing to the target logical qubit operations)

The feasibility of running a utility-generating application depends on these factors. For each problem space, there will be a target run time specific to that problem and that industry, which may be in minutes, hours, or days. While faster quantum operations are desirable to reach a target run time, they are not the end of the story. First, additional qubits can be used to accelerate the application algorithm, mitigating the effects of potentially slow quantum operations. Second, high-fidelity long-range connectivity between qubits can enable vastly more efficient QEC codes which would lower the physical qubit overhead as well as increase the speed with which an algorithm can be executed.

It may come as a surprise, but a low-error rate quantum architecture with long-range connectivity that has comparatively slow quantum operations can compute a problem quicker than a quantum computer with much

faster quantum operations that only has access to nearest neighbour interactions [10].

Scalability is the lifeline of quantum computing

So, the ask from humanity is clear, and what is required for quantum computing to address that ask is clear as well; we need a fault-tolerant quantum computer as soon as possible. Based on the current understanding, this requires a quantum computer with millions of high-quality qubits. It is this understanding that must inform the actions of the decision-makers from the key areas including governments, investors, end users, and quantum computer developers.

It is not obvious at all which of the many entities building quantum computers have the technology capable of reaching the fault-tolerant quantum computing era. It is, therefore, paramount to consider what the key enablers are that we can evaluate to help answer that critical question.

It is one thing to have the best performance in terms of qubit number and error rate at the currently small system sizes, and completely another thing to do it at scale. To better understand the ability of a quantum computer architecture to truly scale up, the following non-exhaustive list of questions may be helpful. A credible plan to build a million-qubit scale quantum computer will address them all:

- How robust and controllable is the qubit at scale? Not all architectures reach the same error rates of operations. An understanding of where those currently stand and can be improved is important. It also appears to be a current occurrence

that as more qubits are being added to a quantum computer, the error rate per qubit increases; i.e., we usually find that for a given architecture, the lower error rates are achieved in quantum computers that have fewer qubits. Decoupling the error rate per qubit from the number of qubits in the quantum computer is crucial. There are two areas of particular importance—(1) the crosstalk when controlling individual qubits (this needs to be removed), and (2) the qubit control system itself. A combination of integrating much of the control system into the chip holding the qubits and a reduction in the correlation between the complexity of the control system and the number of qubits is crucial.

- Is the architecture based on a truly modular approach and, if so, can a sufficiently low error rate and high-speed quantum connection be generated between modules? A vast majority of approaches to quantum computing require a modular approach to be able to scale up all the way. Meaning there is a limitation in the biggest chip that can be produced and, in some cases, the biggest cooling system that can be built, which limits the number of qubits that can be controlled. A modular approach that builds many similar (or maybe identical) modules and connects them together to seamlessly transfer quantum information to obtain a larger system is likely to be the only viable path for a vast majority of architectures. It appears that many of the approaches for building quantum computers do not have a working solution to connect individual modules together. To be clear, without a working solution, it is not possible to scale up.
- Does the required engineering solely rely on commercially available manufacturing solutions? It is crucial to understand and determine if the quantum computing roadmap relies on technologies, material properties, or capabilities that simply have not been invented yet. We know from conventional computing that investments in new chip fabrication and packaging technologies can cost billions of dollars. It is a huge time and cost advantage if currently commercially available manufacturing solutions can be used.
- What error correction techniques can the quantum computer architecture support? The strengths and weaknesses of various quantum hardware approaches determine which quantum error correction codes are feasible. A major factor, which varies greatly across hardware, is the fidelity and rate of long-range interactions between physical qubits. The surface code, relying solely on nearest-neighbour interactions, is a viable choice for most architectures and is particularly favoured in those with limited connectivity, such as superconducting devices. It boasts a high threshold, functioning well with error rates as high as 1% [11]. Reducing error rates below this threshold further lowers the physical qubit overhead needed to achieve a target error rate. Within the broader family of Quantum Low-Density Parity-Check (QLDPC) codes, certain cases require long-range interactions between qubits but are significantly more efficient than the surface code for encoding memory. Only hardware with high-fidelity long-range connectivity can take advantage of these novel QLDPC codes. Additionally, ongoing research in finding efficient means of performing a

universal gate set in QLDPC codes could reduce the overall qubit requirements for achieving quantum advantage in platforms flexible enough to utilize them [12].

- Can the classical computer requirements keep up with the control requirements of the quantum computer? The classical control system of the quantum computer needs to perform various computational tasks while the quantum computer is operating. These tasks must be carried out on the timescale of the clock cycle of a quantum computer and at a fraction of the total time quantum information can be retained for in a qubit (also known as coherence time). The classical control system also needs to scale with the size of the machine. Systems with short coherence times will pose unique challenges to the computational and bandwidth needs of classical computer systems to control those qubits.
- What is the lowest temperature the hardware in the quantum computer needs to be cooled down to? When looking at the temperature question through the lens of requiring an FTQC with millions of qubits, it quickly becomes clear that this is a key question. Most quantum computing architectures require a temperature of a few millikelvins (less than -273 degrees C) to 70 Kelvin. The practical impact of the temperature is that there is orders of magnitude less cooling power available at <1 Kelvin compared to 70 Kelvin. This is important, as operating on qubits adds significant heat to the system that needs to be removed. Sources of heat include heat conducting through the wires connecting the (often at room temperature) control system to the cold qubits as well as heat dissipation. Integration of the control system as close as possible to the qubits is an avenue being pursued; however, note that these complex electronic circuits will need to work at such low temperatures, and that is not a straightforward feat. Simply increasing the cooling power also comes with significant challenges at the lower temperatures. For example, there is currently no available cooling system that could support a million-qubit quantum computer that needs to operate at millikelvin temperatures. This picture changes drastically, however, as the required temperature approaches the 70 Kelvin scale.
- What algorithm speed will the quantum computing architecture ultimately be able to deliver under realistic assumptions? As we showed earlier, it is crucial to understand the overall speed of an algorithm under realistic assumptions that factor in the gate speed, qubit connectivity, and other related factors.
- What is the approximate cost of development and, crucially, what is the expected cost of operation at the million-qubit scale? The current reported pricing of commercially available quantum computers varies greatly but can reach tens of millions of dollars for commercial machines operating with a few tens of qubits. Current initial thoughts across the industry have proposed that the acquisition cost of a one-million-qubit machine needs to be <\$300 per qubit to establish a viable market with a reasonable operating cost that is competitive with high-performance compute centers. It is important to evaluate the required innovations to achieve this across all aspects of the system, with

particular focus on cost-efficient hardware scaling, power-efficient classical control systems, cryogenic and vacuum system needs, and fault-tolerant operation. These areas should be considered from the initial development phase. To arrive at a target development cost of a one-million-qubit system, we can use an often-used 10-100x multiplier which puts the development cost target in the range of \$1B-\$10B.

- Can the supply chain support rapid scalability and is it resilient to unexpected disruptions? Do the supply chain companies need to make massive R&D and capital investments? The supply chain will be critical to the reliable and timely delivery of FTQCs to the market, as can be seen in other advanced technology market segments. Similarly, a challenge will be to transition and qualify new and emerging technologies and their associated manufacturing processes quickly enough to meet market demand. This will be particularly critical as computing power begins to approach utility, as demand would be expected to grow exponentially. This will require early make/buy decisions, establishment of qualified internal infrastructure, and the development of critical supply chains alongside technology development to establish manufacturability, repeatability, and reliability at all levels of hardware and software. Quantum computing architectures that can leverage a commercially available supply chain have an advantage here. Other factors to evaluate include the robustness of the supply chain to external factors. Political instabilities and export control measures need to be protected against as this technology reaches utility and thereby becomes of even stronger national interest.

Closing thoughts

A thought for investors

The venture investment goal in quantum computing is achieving 'quantum advantage'. There will be mid-level winners along the way who potentially sell NISQ machines and generate profits doing so. But these will likely be timing plays for investors, as the need for these systems is only until FTQCs are available. Most of the near-term revenue in quantum computing will go to the big tech companies that allow remote access. Such access will provide a more accessible and less expensive option to customers in the NISQ era, who pay for projects and consulting efforts from trusted sources in order to avoid strategic surprise and to track preparatory steps required for quantum computing in the future. This activity is all important to develop quantum talent, form strategies, and set up industrial processes and standards.

Much of this activity, however, is a distraction from the larger goal of reaching the FTQC scale, which is where the huge value in quantum computing lies.

As discussed above, scaling of the qubit architectures to millions of high-quality qubits is the biggest challenge. Most companies can produce a roadmap that shows the technical milestones required to achieve a step change in scalability. The challenge is that most of these require either 1) research with risk of failure, 2) completed research but unknown manufacturability, or 3) external component improvements to be developed by others. The first and third add significant risk to an investment because

you are essentially making a bet on the competency of a research team and hope that science will not deliver the unexpected surprises it so often does. The second is expensive and lengthy, and you only have limited control over that development. Additionally, semiconductors advanced at the pace they did (which still took decades) due to significant internal research off the profits from existing generation products. Quantum computing will likely not be able to benefit from this type of internal funding, except for big tech companies, because NISQ sales are not significant. Companies with the clearest path towards scalability will likely be the ones to capture R&D budgets to advance their architectures.

A thought for governments

The commercial sector, often driven by the capital markets, has been extremely successful in driving innovation and bringing disruptive technologies to the market. The seeds of those technologies have often resulted from early government support through academic and early-stage commercial R&D support channels such as grants, contracts, and tax relief measures. This union between the government and the private sector in the innovation cycle has been extremely successful. However, it does not work for all sectors/technologies. Quantum computing is one of those, and this appears to be missed by many. We outlined above that to unlock the many applications of quantum computing, to help drive economic prosperity; to ultimately deliver for humanity, we must develop a fault-tolerant quantum computer with millions of qubits. When objectively looking at the timelines for getting to this point, one quickly realizes that it

does not fit the timescales capital markets are accustomed to working with. Expected investor returns, for example, are often on shorter timescales. This has led to a situation where large parts of the current quantum computer sector are being driven towards short-term thinking, focusing on short-term goals that may create some limited economic value on a timescale investors can get behind. This is where some of the hype around NISQ stems from. The real danger here is that it drives the innovators into making decisions that are not aligned with the goal of reaching the fault-tolerant quantum computing scale as quickly as possible.

For governments that are serious about wanting a meaningful outcome in quantum computing, they can no longer rely on their potentially leading economic might to solve this and must shift from short-termism. Without a strong and focused long-term investment strategy, the risks of not harnessing quantum capability are huge. Loss of a technical sector, price points for non-domiciled technology, capital flight to successful regions, industrial and academic decline, and ceding quantum advantage could devastate a nation's industry and economy. But the worst outcome of all would be the deployment of a superior quantum computer against the nation. The question government decision-makers have to ask themselves is: given that the ultimate goal of quantum computing is crystal clear, would a state actor that may currently be inferior economically and militarily allow themselves to be distracted by short-term economic thinking given that this technology could well be their ticket to leapfrogging their nation to economic and defence leadership? The most likely answer should be no surprise to anyone.

With all of the risks clearly on the table, it is surprising that there is still a staggering lack of government support around the world focused specifically on scaling up quantum computing solutions. This is not to say that currently available government support focused on NISQ-type developments are not welcome and important. However, a bifurcated approach is needed, one focusing on supporting NISQ developments and one on driving forward the development of a fault-tolerant quantum computer. Fortunately, the capital markets are nowadays well placed to take care of the NISQ-type developments. The latter is where government support must now focus on. It will be expensive and take time, but taking no action and leaving it to the usual market forces is not an option.

Finally, the ever-increasing export control hurdles for quantum computing coming into effect around the globe pose a significant challenge to the commercial quantum computing sector and risk unintentionally inhibiting the necessary investment into quantum computing development. While protecting national interests is important, it is equally important for governments to enable the development of FTQCs. A proportionate approach to export control restrictions can and must be found that ensures both. ■

References

1. Webber, M., Elfving, V., Weidt, S., & Hensinger, W. K. (2022). The Impact of Hardware Specifications on Reaching Quantum Advantage in the Fault Tolerant Regime. *AVS Quantum Science*, 4(1), 013801. DOI: 10.1116/5.0073075.
2. Gidney, C., & Ekerå, M. (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5, 433. DOI: 10.22331/q-2021-04-15-433
3. Alexander M. Dalzell, Sam McArdle, Mario Berta, Przemyslaw Bienias, Chi-Fang Chen, András Gilyén, Connor T. Hann, Michael J. Kastoryano, Emil T. Khabiboulline, Aleksander Kubica, Grant Salton, Samson Wang, and Fernando G. S. L. Brandão, "Quantum Algorithms: A Survey of Applications and End-to-End Complexities," *Quantum Information at RWTH Aachen University*, 2023.
4. Jin-Yi Cai, "Shor's Algorithm Does Not Factor Large Integers in the Presence of Noise," *Science China Information Sciences*, vol. 67, no. 7, Article 173501, June 2024.
5. Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H. Booth, and Jonathan Tennyson, "The Variational Quantum Eigensolver: A Review of Methods and Best Practices," *Physics Reports*, vol. 986, pp. 1-128, 2022. DOI: 10.1016/j.physrep.2022.08.003.
6. Lee, J., Berry, D. W., Gidney, C., Huggins, W. J., McClean, J. R., Wiebe, N., & Babbush, R. (2021). Even More Efficient Quantum Computations of Chemistry Through Tensor Hypercontraction. *PRX Quantum*, 2(3), 030305. DOI: 10.1103/PRXQuantum.2.030305.
7. Santagati, R., Aspuru-Guzik, A., Babbush, R., Degroote, M., González, L., Kyoseva, E., Moll, N., Oppel, M., Parrish, R. M., Rubin, N. C., Streif, M., Tautermann, C. S., Weiss, H., Wiebe, N., & Utschig-Utschig, C. (2024). Drug Design on Quantum Computers. *Nature Physics*, 20(4). DOI: 10.1038/s41567-024-02411-5.
8. Yuri Alexeev, et al., "Quantum-centric Supercomputing for Materials Science: A Perspective on Challenges and Future Directions," *Future Generation Computer Systems*, 2024, pp. 577-597. DOI: 10.1016/j.future.2024.04.015.
9. Chakrabarti, S., Krishnakumar, R., Mazzola, G., Stamatopoulos, N., Woerner, S., & Zeng, W. J. (2021). A Threshold for Quantum Advantage in Derivative Pricing. *Quantum*, 5, 463. DOI: 10.22331/q-2021-12-15-463.
10. Wan, K., Webber, M., Fowler, A. G., & Hensinger, W. K. (2024). An iterative transversal CNOT decoder. *ArXiv preprint*, arXiv:2407.20976. DOI: 10.48550/arXiv.2407.20976
11. Oscar Higgott and Craig Gidney, "Sparse Blossom: Correcting a Million Errors per Core Second with Minimum-Weight Matching," *Physical Review X*, vol. 13, no. 3, 2023. DOI: 10.1103/PhysRevX.13.031007.
12. Guanyu Zhu, Shehryar Sikander, Elia Portnoy, Andrew W. Cross, and Benjamin J. Brown, "Non-Clifford and Parallelizable Fault-Tolerant Logical Gates on Constant and Almost-Constant Rate Homological Quantum LDPC Codes via Higher Symmetries," *arXiv:2310.16982 [quant-ph]*, 2023.

Cover photo by Yulian Alexeyev on Unsplash



About Universal Quantum

Universal Quantum is pioneering scalable trapped ion quantum computers, poised to revolutionise multiple industries.

With breakthrough technologies like UQConnect, enabling world-record quantum connections between chips, and UQLogic, offering robust and scalable qubit control, we are paving the way to utility-scale quantum systems. Crucially, our innovative quantum computers are being built using readily available manufacturing technologies and operate at a practical temperature of 70K.

Contact us at info@universalquantum.com to find out more about Universal Quantum..

www.universalquantum.com