Publication date:
December 2023
Author:
Sam Lucero

In Pursuit of Fault Tolerant Quantum Computing

Brought to you by Informa Tech

Commissioned by: Microsoft

Contents

The Path to Quantum Computing Today

Quantum computing (QC) may one day enable solutions to perform computational challenges that classical computing will never find traction on – at least not on any human-relevant timescale. While this would almost certainly lead to significant downstream scientific and technological breakthroughs, including progress on responding to climate change and finding key new medicines for some of humanity's greatest health challenges, achieving this kind of "quantum advantage" is both far from certain, and difficult to measure as a series of interim goals. **In this whitepaper, Omdia discusses why clarity on progress towards quantum advantage is hard to find, why this lack of clarity is a challenge for the industry, and how the industry might find benefit using a framework to think about progress towards quantum advantage.**

Quantum advantage is a current term of art in the industry to denote QCs' ability to outperform classical computers (having mostly replaced the older, and now mostly disfavored, phrase "quantum supremacy"). Unfortunately, "quantum advantage" is used in different ways and in different contexts, leading to confusion by potential and current adopters about the current and possible future capabilities of QCs, and what adopters can expect in terms of timelines. This confusion increases uncertainty and risk regarding the investments that adopters should make in QC, and when. Should adopters start to invest now, even if only to experiment and learn about QC technology? What scale of investment would lead to a useful outcome?

Quantum computing (QC) may one day enable solutions to perform computational challenges that classical computing will never find traction on – at least not on any human-relevant timescale.

Tracking the Early Adopters

Omdia's research suggests that adoption of QC technology is already relatively strong, at least among government, academic, and large commercial users (think "Global 2000" companies). We estimate there are nearly 500 adopters of QC technology in 2023, and this will rise to nearly 6,300 by 2027 and to over 18,000 by 2032. How much are these adopters investing in their efforts? Based on anecdotal as well as survey-based findings, the typical project spend is in the low single digit millions of dollars per year. For example, in Omdia's 2023 QC adopter survey, the largest segment (26%) of respondents in China, Germany, and the US stated that their organizations had committed to an annual budget

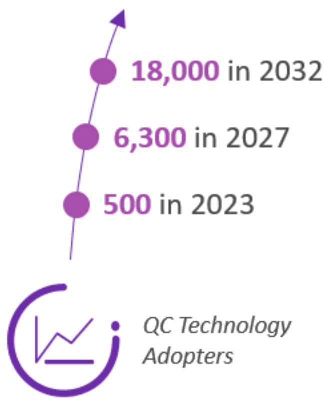

of between $1 million and $2 million. And by the way, 13% of respondents selected the response of "More than $5 million".

What are these adopters getting for their investments? In most cases, these seem to be efforts to experiment with QC technology to learn about its capabilities, upskill staff and develop internal QC expertise, and, in some cases, to start to develop intellectual property rights (IPR) protections related to quantum algorithms adapted for specific applications.

For example, Omdia knows of one major life sciences company that spent what Omdia estimates as about $2.5 million for 500 hours of quantum processing unit (QPU) time to run a benchmarking project. The company did not develop a directly operational capability with this investment – rather, it plans to run this benchmark again in a few years to determine how state of the art for QC technology has advanced for the company's specific needs.

Spending millions of dollars on experiments and benchmarking exercises is not an option for all potential adopters. However, many adopters view such investments as an "insurance policy" against the very real possibility that large scale, fully fault tolerant QCs (FTQCs) could be transformative across many industries where complex computational challenges loom large, such as life sciences, manufacturing, energy, or chemicals and materials. No company wants to be at a disadvantage when, and if, QC technology passes a "quantum advantage" inflection point (if nothing else, the recent and sudden prominence of generative AI and the resulting scramble among executives to determine their "ChatGPT play" has been an object lesson on the need to prepare to run the race before the race starts).

In some limited cases, however, adopters do claim to find an "advantage" today, even with current noisy, intermediate-scale quantum (NISQ) QCs: in the adopter survey referenced above, 29% of respondents stated that their organization already sees a "commercially relevant" advantage to using QCs (this was particularly true for respondents of Chinese-headquartered organizations, interestingly). This is a surprising result, and one that drives to the heart of the challenge facing in the industry of just how we determine when we have a "quantum advantage". Leaving aside a (probably very large) portion of these respondents who may define a "commercially relevant" advantage as the ability simply to experiment with QC technology and to "prepare to run the race", we find that addressing the definitional challenge is key, as we now discuss below.

No company wants to be at a disadvantage when, and if, and when QC technology passes a "quantum advantage" inflection

Breaking Down Quantum Advantage

Omdia believes that "quantum advantage" can be decomposed into three consecutive and increasing levels of capabilities relative to classical computing:

- **Commercial** advantage is in the eye of the beholder, when an adopter believes they receive some commercial benefit to using QC compared to the traditional classical computing solution they would normally use. This type of advantage may be in speed but could also be measured in the quality of the results or the cost of achieving the results.

For example, one air cargo transport company conducted a test using historical logistical data in which a QC was able to show how to successfully load and transport the 20% of cargo in the dataset that the classical logistics system had been unable to allot space for on the planes. Had the QC been used operationally, the ability to move this cargo in the first instance would have resulted in more efficient revenue generation since the cargo company is only paid for cargo that reaches its destination.

A key point to emphasize here is the comparison of the QC results to the "traditional classical computing solution" that would normally be used for the application. One could argue that the adopter should bypass NISQ era QC and move directly to developing a classical high-performance computing (HPC)-based solution instead. Nevertheless, organizations can have reasons to prefer trying QC over trying classical supercomputers (such as a desire for data privacy combined with an unwillingness to build and operate a supercomputer).

- **Computational** advantage is more objective but still, in many or most cases, empirical rather than formally provable or absolute. For instance, in June 2023, IBM Quantum announced that it had achieved a more accurate computational result for a problem of commercial interest than could be returned from a classical supercomputer (i.e., not just "the traditional alternative solution" but the best that classical computing can offer). Interestingly, within a week several other research groups responded that they had found better classical algorithms that outperformed the IBM Quantum result. IBM Quantum positions this as a healthy back and forth between the QC and classical computing camps. This result is important, however, because it highlights that a strong rationale for using QCs operationally may exist even in the NISQ era, if, as expected, QC technology continues to advance along a spectrum of growing scale and increasingly performant error suppression and error mitigation capabilities. If these capabilities advance sufficiently, NISQ-era QCs may be able to provide a computational advantage in some instances over classical computing. It may be unclear how long each period of QC computational advantage might last before QCs are once more surpassed by classical HPCs, but we might expect that over time, and with further QC technology advancements, these periods grow increasingly longer. This would be highly positive in moving customers from small-scale experiments to large-scale operations. However, this conjecture is very speculative, and we need to stress that many experts in the field believe that NISQ-era QCs will never provide a computation advantage over classical computing.

- **Tractability** advantage is what many people think of as "quantum advantage" — the point at which no classical supercomputer could conceivably match the results of a QC, at least on human-relevant timescales (in other words, if it took a supercomputer 100 years, or 1,000 years, to complete a calculation, it's probably useless as a practical matter). It will almost certainly take an FTQC to achieve this, and most estimates are that this will require another 10 or more years to develop. There is also the consideration of where the industry will see a true superpolynomial or exponential advantage, as we will discuss further below. Although, even if tractability advantage is "limited" to physical simulation, this alone would be a world-changing result with vast economic and societal impacts.

Quantum Tractability Advantage and Microsoft's Goal

Simply put, while isolated examples of "quantum commercial advantage" are periodically announced (almost always in the context of experimentation rather than operational use), these are very far from the full potential of QC if FTQCs are achievable. Commercial advantage might provide an incremental benefit on a case-by-case basis, but will not be the inflection needed to, for example, meet Microsoft's aspirational goal to "compress the next 250 years of scientific discovery into the next 25." For that, the industry will need to achieve "quantum tractability advantage".

Achieving such a tractability advantage will require perhaps up to ten or more years and necessitate significant scientific and technical breakthroughs. These advances will be necessary in both the development of FTQCs (i.e., large volumes of logical qubits) and, if possible, the discovery of new superpolynomial quantum algorithms. On the hardware front, there are half a dozen promising physical qubit modalities, ranging from various types of superconducting circuit qubits to nitrogen vacancy (NV) centers.

Notably, Microsoft Azure Quantum recently passed a key milestone in developing its own type of qubit—a topological qubit—that should be significantly more robust than current, more established qubit technologies, if fully realizable.

To give a sense of the potential performance benefit of topological qubits, Microsoft Azure Quantum is targeting error rates of 10^-4 (compared to current state of the art error rates ranging from 10^-2 to 10^-3) and suggests that a path exists to achieve error rates as low as 10^-6. While quantum error correction codes would still be needed to achieve the 10^-12 to 10^-18 error rates necessary for FTQC, this order(s) of magnitude improvement will have a material benefit in the types of QEC codes that are used, and how they are used, which in turn has a benefit in the ratio of physical to logical qubits required for FTQC. This improvement in ratio scaling, in turn, positively impacts how quickly FTQCs can progress along the "quantum supercomputer" performance spectrum.

Relatedly, Microsoft expects their topological qubits to lead to benefits in QPU size, per operation speed, and relieving the I/O bottleneck between the qubits and the classical control system. More specifically, Microsoft plans to fit more than a million physical topological qubits on a wafer that is smaller than the security chip on a credit card. This should negate the need for complex, entangled quantum communications between multiple modules in a QPU. Likewise, Microsoft's goal is for a per operation speed as low as 100 nanoseconds, on par with superconducting circuit qubits. Finally, Microsoft has developed a cryoCMOS approach to control chip design to enable placement of control circuitry physically close to the qubits, in the same cryogenic system, albeit in a different, thermally isolated segment. It is critical that lower physical error rates do not come at the price of reduced speed in the case of topological qubits. Moreover, they also have a path to solving the I/O bottleneck through cryogenic CMOS control because their operations are fundamentally digital, not analog.

The Industry Challenges Ahead

The major project for QC hardware will be to increase both the number of qubits and to reduce the error rates of the physical qubits, which is directly implicated in how effectively quantum error correction (QEC) codes can be implemented in the creation of logical qubits. Essentially, the physical qubit error rate must be about 100 to 1,000 times lower than the QEC code breakeven error threshold to enable logical qubits based on a “practical” ratio of physical to logical qubits in the range of several hundred to one or a thousand to one. (If the physical qubit error rate is just slightly below the QEC code breakeven error threshold, each logical qubit would require a completely unrealistic several hundred thousand physical qubits).

The need for superpolynomial quantum algorithms is an underappreciated challenge in the industry. Basically, due to the overhead inherent in QC operation compared to classical computers, a quantum tractability advantage is likely only achievable (even for FTQCs) if the quantum algorithm offers a superpolynomial, and ideally an exponential, advantage over the classical analogue. Unfortunately, to date it appears that to achieve such a speedup requires an underlying structure to the computation itself (rather than be a “black box” computation requiring each potential solution to be tried in parallel). Shor’s Algorithm for factorization leverages such structure. Likewise, algorithms for the simulation of quantum mechanical physical systems show such structure (i.e. the way electrons interact with nearest neighbors), although, the very fact that quantum systems are governed by quantum mechanics inherently make them exponentially harder to model and compute on classical computers compared with using a QC.

However, we haven’t yet found similar structure for algorithms in the quantum machine learning (QML) and combinatorial optimization use cases, at least in a broad-based way applicable to many or most QML and combinatorial optimization algorithms. Perhaps such structure will be found eventually with the efforts of the growing quantum community to unlock new solutions. If not, it’s very likely that quantum tractability advantage will be focused on applications requiring the simulation of quantum mechanical physical systems like chemistry and materials science. Even if this turns out to be the case, this would arguably be a fantastic, earth-changing result, leading to the types of advances mentioned at the outset of this note.

Even if quantum tractability advantage is limited to physical simulation, this would arguably be a fantastic, earth-changing result, leading to the types of advances mentioned

Nevertheless, the semantic ambiguity around the term “quantum advantage”, the diversity of approaches to instantiating physical (and logical) qubits, and the complexity in finding superpolynomial quantum algorithms, all make measuring and understanding the industry’s progress towards quantum tractability advantage exceedingly hard to do. Particularly for QML and combinatorial optimization use cases, we may only be able to understand advantage as an empirical result, requiring many instances of such an advantage being announced by multiple adopters before the industry can collectively have confidence that a specific implementation of QC technology truly does offer a benefit over classical computing alternatives. This both makes it challenging for potential adopters to organize their investments in QC technology, and, concurrently, for vendors to find sales traction.

Omdia notes that some QC software vendors are emphasizing 'quantum-inspired' functionalities, or leveraging today's supercomputing and artificial intelligence capabilities, recognizing the need to deliver value in the short-term while the QC market is still developing.

Adopters need benchmarks and a framework for thinking about the development of QC technological capabilities to rationally allocate investments. Originally (and still), the total number of physical qubits has been one simple benchmark for QC functionality, and most vendors announce the number of qubits offered by their system. More recently, “Quantum Volume” (a measure that seeks to embed multiple functional parameters in one number) has gained popularity. However, Quantum Volume, and related measures that focus on the performance of physical qubits, will offer an incomplete view of the performance of FTQCs as we move into a regime of computation using logical qubits. A broader view that extends into the FTQC era will be needed.

rQOPS: Measuring FTQC Computational Performance

Microsoft’s proposed “reliable quantum operations per second” (rQOPS) metric is a new standardized (i.e., applicable across different vendors’ QC systems) performance benchmarking framework. Standardized benchmarking tools are useful because QCs are highly complex and variable in their capabilities and attributes based on their qubit type and vendor of origin. Variation comprises not only number of qubits and qubit error rates, but also factors such as qubit-to-qubit connectivity and available native gate sets. As mentioned above, not every adopter will have the time or resources to conduct customized benchmarking projects focusing on their specific use case of interest. Therefore, the QC vendor community has taken it upon itself to offer benchmarking approaches that are as simple to use as possible, while still enabling reasonable comparisons between systems.

Several QC benchmarking approaches exist today. In Omdia’s view, the most popular and widely used of these approaches is IBM Quantum’s proposal for a “Quantum Volume” (QV) metric. It will be useful to compare QV and rQOPS to illuminate those cases in which rQOPS is meant to provide benefits that aren’t offered by current approaches.

The rQOPS metric has been defined with the purpose of measuring the system's performance and helping us understand its ability to solve impactful problems. As it captures the capabilities of a quantum supercomputer, it is calculated by multiplying the volume of logical qubits in a QC system by the QC hardware’s logical clock speed and is expressed with a corresponding logical error rate giving the maximum allowable error rate of the system’s operations on the logical qubits. By measuring how many reliable quantum operations can be executed in a second, this metric enables measurement of the scale, speed, and reliability of an FTQC.

As a formula, rQOPS is given by the number Q of logical qubits in the quantum system multiplied by the hardware's logical clock speed f:

rQOPS = Q· f

It is expressed with a corresponding logical error rate pL, which indicates the maximum tolerable error rate of the operations on the logical qubits.

In contrast, QV was designed by IBM Quantum explicitly to measure the performance of "near-term devices with a modest number of qubits". QV compares the ability of different NISQ-era QCs to perform a relatively small-scale quantum computation directly on physical qubits. The QV metric does this by encapsulating in one number a complex evaluation of a QC's number of physical qubits, qubit error rate, qubit-to-qubit connectivity, and the available native gate set. A key factor that anchors QV squarely in the NISQ era is the use of the calculated difference between the results of a random circuit computation performed on the QC hardware of interest versus the same computation run on an ideal simulator, as one of the inputs in the overall QV metric.

Just recently (November 2023), IBM Quantum introduced a new benchmarking proposal it calls "Layer Fidelity" (LF). LF is not meant to replace QV, but rather to extend the ability to benchmark QCs in ways that QV is unable to do. The LF benchmarking method identifies some advantages such as increased scale, full system measurement, system and component-level measurements, continuous and flexible measurement. It's presented as making it more relevant for evaluating quantum computers in the context of error mitigation.

It is important to note, however, that LF is still a measure of physical qubits, albeit at a scale larger than is emulatable by classical computers. Also, LF is still not a benchmark of speed. To Omdia's knowledge, circuit layer operations pers second (CLOPS) is the only speed benchmark used currently, apart from Microsoft Quantum Azure's rQOPS proposal. Other benchmarks measure quality and scale, and this is true for LF as it is for QV.

In short, rQOPS is analogous to the floating-point operations per second (FLOPS) metric used to compare classical supercomputers: it provides a comparison of the speed with which different FTQCs can perform a given fault-tolerant circuit. This speed comparison is critical to know because it looks likely that fault-tolerant circuits (e.g., Shor's Algorithm) will require many billions, or even trillions, of operations (i.e., gates) to compute. Even using an FTQC, the per operation speed becomes a potential limiting factor, given that two different QCs could, for example, exhibit several orders of magnitude of difference in speed, potentially rendering a computation practical on the faster QC and impractical on the slower QC. The figure below shows Microsoft's estimate of the speedup needed to move, for example, from computing chemistry simulations on practical time scales to computing material science simulations practically.

Figure 1. Illustration of rQOPS at work

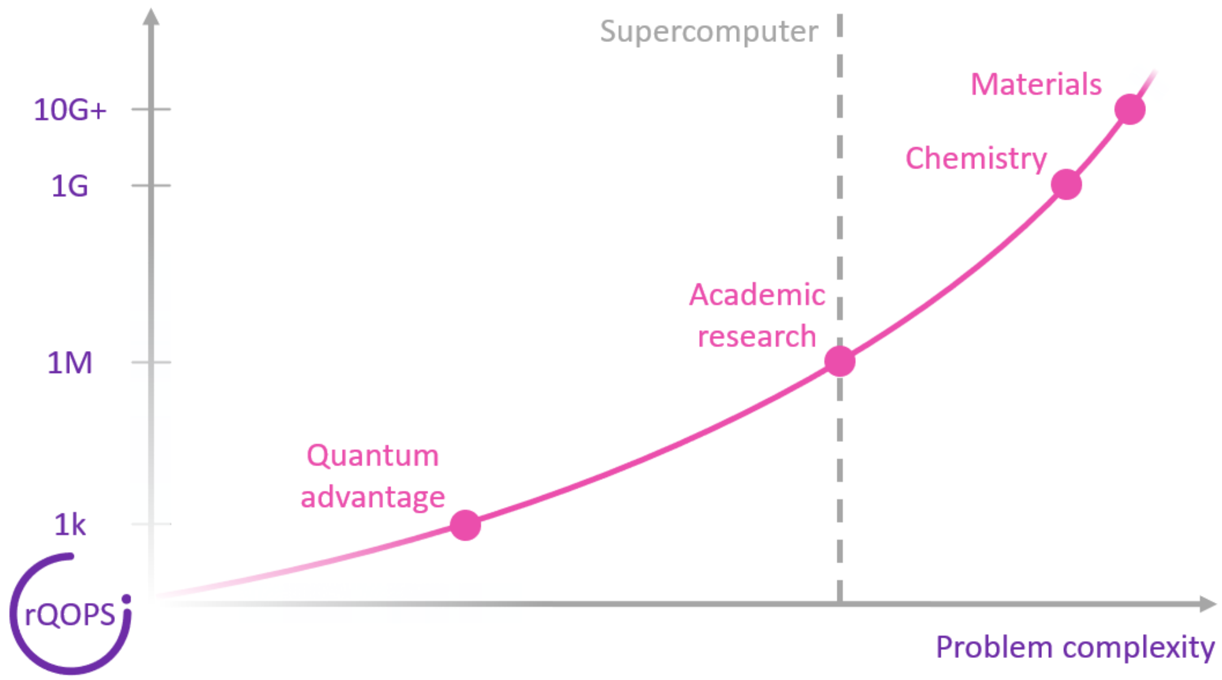

Source: Microsoft, Omdia

Introducing the Quantum Computing Implementation Levels Framework

The Quantum Computing Implementation Levels framework is a useful approach for thinking about QC technological progress beyond the NISQ era into the fault tolerant regime. The framework is meant to work synergistically with the rQOPS metric proposal outlined above. The framework contains three sequential levels:

Level 1—Foundational: Quantum systems that run on noisy physical qubits, including all of today's Noisy Intermediate Scale Quantum (NISQ) computers. Level 1 denotes the current state of QC technological development, which relies on small scale computation directly on physical qubits. This level aims to provide a foundation to start experimenting with quantum technologies. The industry is only starting to experimentally show that logical qubit error rates can fall below physical qubit error rates. There is still some doubt, particularly in a small share of the academic physics community, that the industry can successfully deliver FTQC. Omdia believes the probability of ultimately passing out of Level 1 to achieve working logical qubits with error rates below physical qubits is high; we estimate the probability at above 90%.

2

Level 2—Resilient: Quantum systems that operate on reliable logical qubits. Level 2 is the successful development of effective logical qubits. At this stage, QCs should enable experimentation with quantum algorithms written for logical qubits, and these experiments may demonstrate increasingly robust commercial advantages to using QCs rather than classical computing. Again, the advantage may be incremental and empirical, rather than formally provable in a complexity theoretic manner, but there could be growing certainty in the industry that operational benefits will follow from investments into QC. At this stage, there will still be many different types of underlying physical qubits and using different modalities may be akin to working with GPUs as opposed to CPUs. Different qubit types will have different connectivity models, modularity considerations, physical error rates, and coherence times, for example.

3

Level 3—Scale: Quantum supercomputers that can solve impactful problems which even the most powerful classical supercomputers cannot. Level 3 is solidly in the "quantum tractability advantage" regime mentioned above. At this level, the rQOPS metric becomes central to capturing what a "quantum supercomputer" can achieve in solving the types of exponentially challenging computations described at the outset. Interestingly, Microsoft suggests that at least a million rQOPS—a million reliable operations per second—will be needed by a QC able to power exponential class quantum algorithms, with an associated logical error rate of at most 10^-12 (or one only one error for every trillion operations) for simple quantum materials simulation. The company believes one billion rQOPS, with a logical error rate of at most 10^-18, will be needed for complex chemical and materials science simulation. (Microsoft strongly believes that exponential class quantum algorithms for QML or combinatorial optimization will not be found, hence their focus on physical simulation, though this view is not universally shared in the industry).

While QC holds vast potential for helping to solve some of society's most significant challenges, the technology and market are still at an extremely early and ill-defined stage of development, making it challenging for current and potential adopters to plan their investments effectively. Frameworks to think about "quantum advantage" and figures of merit to measure progress towards FTQC will be enormously helpful in increasing the efficiency of activity in the industry. The **Quantum Computing Implementation Levels** and **rQOPS** figure of merit, proposed by Microsoft, are useful contributions to this effort at organizing thinking about such progress.

Frameworks to think about "quantum advantage" and figures of merit to measure progress towards FTQC will be enormously helpful in increasing the efficiency of activity in the industry.

Finding Quantum Advantage: Simulating Quantum Mechanical Physical Phenomena

Given the preceding discussion, and leaving aside notions of “commercial advantage”, where might the industry find the most success in achieving a quantum “tractability advantage”—that is, a true quantum speedup over the best possible classical computing solutions? The most conservative, and likeliest, answer lies with those applications for which a quantum algorithm exists that provides at least a superpolynomial—not simply quadratic—speedup.

The basic reason for this likelihood is that quantum computers bear the burden of an inherent “overhead” relative to classical computers. While physical qubit per operation speeds vary widely between different qubit modalities, with fast times for superconducting circuit qubits, and much slower times for trapped ion qubits, for example, in general logical qubit per operation speeds are on the order of a million times slower than classical compute operations.

This per operation slowdown exacerbates a corollary input/output (I/O) overhead burden: for the foreseeable future, a “large” FTQC may only encompass about 10,000 logical qubits. This is a far cry from a classical graphics processing unit (GPU) with billions of transistors. And classical data can only be loaded one datapoint at a time into each of these 10,000 qubits. So, if the FTQC starts with a million times fewer processing units (qubits) compared to a hypothetical GPU with 10 billion transistors, and each of these processing units is a million times slower than the GPU’s transistors, then we face a constant, massive slowdown in the speed with which we can load our classical data, process the data, and readout the data into classical form. (This hypothetical example about I/O avoids discussion of the idea of using quantum random access memory—QRAM—which does not exist currently and seems unlikely to exist any time soon.)

> *The overhead described implies that FTQCs will offer a speedup advantage when the problem involves a “small data, big compute” computation and we have at least a super-polynomial quantum algorithm to leverage at the outset of this note.*

The overhead described above implies that FTQCs will offer a speedup advantage when the problem involves a “small data, big compute” computation and we have at least a superpolynomial quantum algorithm to leverage:

- **Small data:** We only want to load a relatively small amount of data into the FTQC, to avoid the overhead mentioned above. While this generally would preclude typical “big data” endeavors, such as loading a large ML training dataset to find patterns, there may still be clever ways to statistically reduce the size of a classical dataset such that a much smaller, albeit still statistically representational subset of the data can be used successfully on an FTQC. Mostly, though, we will want to limit ourselves to highly structured problems where a complex problem statement can be described very concisely.

- **Big compute:** We want to focus on problems that are combinatorially complex; where the solution becomes exponentially more difficult to compute with each new variable or restraint added. This problem profile has the best possibility for combining “small data” with a so-called “combinatorial explosion” in compute complexity that is likely to overwhelm a classical computer while still being tractable to an FTQC.

- **Superpolynomial quantum algorithm:** Because of the constant overhead limitations on FTQCs mentioned above, we want quantum algorithms that solve problems in at least superpolynomially fewer, and ideally exponentially fewer, “oracle calls” compared to a classical computer. This is a subjective consideration: even a quadratic quantum algorithm (that solves problems with several oracle calls that is the square root of the number made by a classical computer) will, eventually, overcome this overhead and produce a quicker answer. However, if “quicker” still means getting the answer “in months or years”, then that may not be very practical for our purposes. Furthermore, some problems may simply not be “large enough” to eventually reach the point where a quadratic speedup kicks in; in such a case, the classical computer simply solves the problem faster than the overhead challenged FTQC. Indeed, Omdia has spoken with practitioners who complain of needing to artificially “expand” an optimization problem (for example, by loosening hard restraints) before they start to see a quantum speedup in their experiments.

To date, the types of problems that best fit the characteristics listed above generally fall in the areas of simulating quantum mechanical systems for chemistry and materials science use cases, along with factoring large prime numbers. Linear algebra problems also see an exponential speedup, although the I/O overhead burden means that these problems need to be of the highly structured, “small data” variety (e.g., “physics-inspired”) to be practical for quantum computation.

Quadratic speedups exist for other types of quantum algorithms, such as with Grover’s Search algorithm, but these generally don’t get us over the hurdles described above. Again, our focus here is on quantum speedups, not other forms of “advantage”. For example, we’re not considering potential benefits provided by QML that are not speedups, but rather gains in model expressability or generalizability.

It’s important to note that the current lack of superpolynomial quantum algorithms in use cases outside of physical simulation and large prime number factoring don’t preclude such algorithms from being found in the future. For example, in the video interview with Microsoft’s Mattias Troyer referenced in the links in the Appendix, it’s noted that one of Microsoft’s first quantum algorithms for physical simulation was “a billion times slower” than a later algorithm developed by the company. It’s conceivable that future advancements will be made in other areas, and, indeed, quantum algorithm development is an active area of research.

However, we must also note that some QC researchers, such as Scott Aaronson at UT Austin, believe that superpolynomial speedups will only be found in cases where the quantum algorithm can exploit some inherent structure of the problem. This raises the bar of difficulty for finding superpolynomial quantum algorithms for a broad set of problems.

Figure 2. Three possible scenarios for QC technology and market development

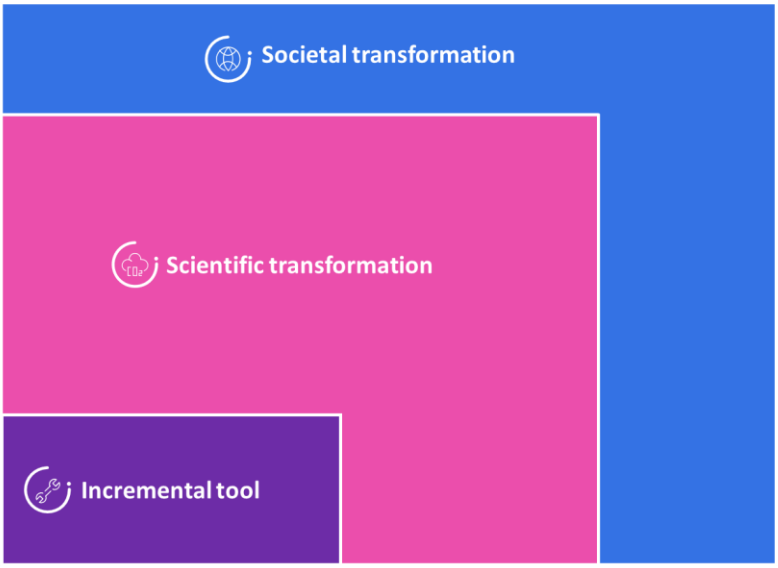

Source: Omdia

Synthesizing the points above leads to three possible scenarios for how the QC market will develop in the coming years and impact on the larger economy (which we illustrate above in Figure 2). In the first scenario, “Incremental tool”, FTQCs prove untenable, and we’re limited to narrow instances of commercial advantage using QC computations running directly on increasingly robust physical qubits. In this scenario, there is a benefit to QC, but it is limited in scope and generally seen on a case-by-case basis. The advantage in this case will be subjective and empirical in nature.

The second scenario, “Scientific transformation”, mostly closely aligns with Microsoft Azure Quantum’s aspirational goal of condensing the next 250 years of chemistry and materials science research into the next 25 years. It holds that FTQCs do become available and are paired with existing superpolynomial quantum algorithms in the areas of quantum mechanical physical simulation to make existing scientific advances that prove transformative to our current scientific processes.

Finally, the third scenario, “Societal transformation”, expresses the maximally optimistic view that not only do FTQCs become available, but that key advances are made in areas such as finding a broad-based set of superpolynomial and exponential quantum algorithms beyond just physical simulation and prime number factorization, QRAM becomes practical, and, perhaps, new techniques are developed to greatly reduce the QC overhead burdens described above.

Appendix

About Microsoft

Microsoft (Nasdaq "MSFT" @microsoft) enables digital transformation for the era of an intelligent cloud and an intelligent edge. Its mission is to empower every person and every organization on the planet to achieve more.

Further reading

Quantum Advantage: Hope and Hype - Microsoft Azure Quantum Blog

Evaluating Quantum Computing Research | LinkedIn

Quantum Computing: The Importance of Fidelity and Error Correction | LinkedIn

Measuring Quantum Volume: Qiskit Online Course

Disentangling Hype from Practicality: On Realistically Achieving Quantum Advantage | Communications of the ACM

How can we solve quantum problems today and in the future? | Microsoft

Author

Sam Lucero

Chief Analyst, Quantum Computing

Applied Intelligence

Sam.Lucero@omdia.com

Get in touch

www.omdia.com
askananalyst@omdia.com

Omdia consulting

Omdia is a market-leading data, research, and consulting business focused on helping digital service providers, technology companies, and enterprise decision-makers thrive in the connected digital economy. Through our global base of analysts, we offer expert analysis and strategic insight across the IT, telecoms, and media industries.

We create business advantage for our customers by providing actionable insight to support business planning, product development, and go-to-market initiatives.

Our unique combination of authoritative data, market analysis, and vertical industry expertise is designed to empower decision-making, helping our clients profit from new technologies and capitalize on evolving business models.

Omdia is part of Informa Tech, a B2B information services business serving the technology, media, and telecoms sector. The Informa group is listed on the London Stock Exchange.

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help your company identify future trends and opportunities.