

QIR

Quantum
Index Report
2025

MIT INITIATIVE ON THE DIGITAL ECONOMY

M
DIGI
TAL

► How to cite this report

Ruane, J., Kiesow, E., Galatsanos, J., Dukatz, C., Blomquist, E., Shukla, P., “The Quantum Index Report 2025”, MIT Initiative on the Digital Economy, Massachusetts Institute of Technology, Cambridge, MA, May 2025.

The Quantum Index Report 2025 by Massachusetts Institute of Technology is licensed under **CC BY-ND 4.0** [Attribution-NoDerivatives 4.0 International](https://creativecommons.org/licenses/by-nd/4.0/).

► Interactive website and public data

The Quantum Index Report 2025 is accompanied with interactive tools available on our website (qir.mit.edu) and we share our raw data with the community available to download from our website (qir.mit.edu/data).

In memory of Shawneric Hachey, whose unique talent and dedication shaped the way this project is presented today.



Center for Quantum Networks
A National Science Foundation Engineering Research Center



This work was supported by the Engineering Research Centers Program of the National Science Foundation under NSF Cooperative Agreement No. 1941583. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the National Science Foundation.

▶ Team

Jonathan Ruane, Principal Investigator and Editor-in-Chief
MIT Sloan School of Management
MIT Initiative on the Digital Economy

Elif Kiesow, Senior Researcher and Project Manager
MIT Initiative on the Digital Economy

Johannes Galatsanos, Researcher
MIT Initiative on the Digital Economy

Carl Dukatz
Accenture

Edward Blomquist
Accenture

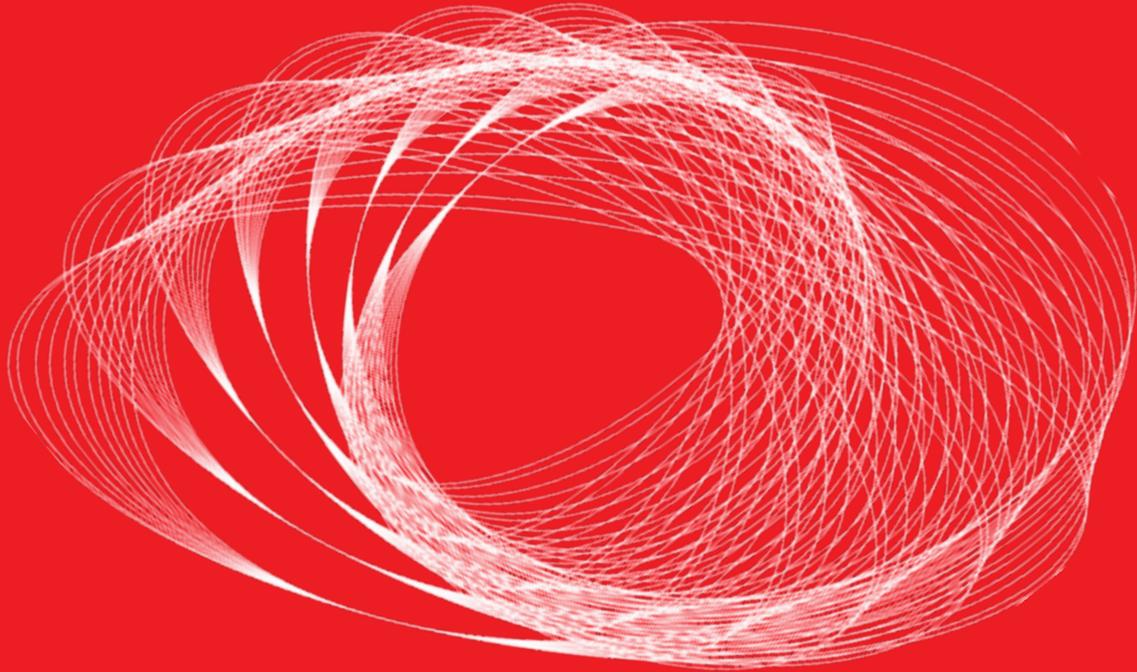
Prashant Shukla
Accenture



This research is a collaboration between Accenture and the MIT Initiative on the Digital Economy (IDE) and was performed under the MIT and Accenture Convergence Initiative for Industry and Technology.

Table of Contents | 2025 Quantum Index Report

Key Insights	<i><u>page 7</u></i>
Executive Summary	<i><u>page 9</u></i>
Chapter 1 Patents	<i><u>page 17</u></i>
Chapter 2 Academic Research	<i><u>page 24</u></i>
Chapter 3 Venture Funding	<i><u>page 41</u></i>
Chapter 4 Quantum in Corporate Communications	<i><u>page 47</u></i>
Chapter 5 Policy	<i><u>page 53</u></i>
Chapter 6 Workforce	<i><u>page 65</u></i>
Chapter 7 Education	<i><u>page 71</u></i>
Chapter 8 Public Opinion	<i><u>page 77</u></i>
Chapter 9 Quantum Networking	<i><u>page 89</u></i>
Chapter 10 Quantum Processor Benchmarking	<i><u>page 93</u></i>
Appendix	<i><u>page 118</u></i>



10 | Quantum Processor Benchmarking

It is challenging for non-experts to easily understand the performance of quantum computing today. Without this understanding, making predictions about investment, commercial deployments, use case testing, and overall strategy is prohibitive. This opacity has many drivers, including the nascent state of the technology, the existence of multiple modalities (types of quantum computers), the lack of independently verified universal performance metrics, and the context-dependent connection between hardware devices and quantum algorithms.

Benchmarks

To enable better insight into the current state of quantum computing performance, we indexed and analyzed published data on over 200 Quantum Processing Units (QPUs) from 17 countries, including retired, prototype, current, and announced QPUs. As of April 2025, there are over 40 commercially available QPUs from at least two dozen manufacturers.¹

It is important to preface that quantum computers remain below the performance capabilities of classical computers for all useful tasks today.

A useful analogy when considering QPU benchmarking is to consider the comparison of racecars. Racecars benefit from large horsepower and torque ratings, lightweight construction, independent suspensions, and performance aerodynamics. However, lap times on a specific racetrack are the best overall measure of how well the various elements perform together. QPUs have similar construction, design and performance idiosyncrasies. There are three primary categories of QPU benchmarks that are useful to consider:

▶ **Physical benchmarks**

e.g., number of qubits and fidelity of the qubit gates.

These are the core metrics of a QPU. They are akin to the weight and torque values of a racecar. While they are objective measures, they only provide a partial insight into the likely overall performance when described in isolation. Physical benchmarks are the category that QPU manufacturers are most likely to disclose and are a focus area for the analysis below.

▶ **Aggregated benchmarks**

e.g., Quantum Volume, CLOPS, and Logical Qubits.

These are various combinations of physical benchmarks. In the car analogy, these benchmarks are similar to power-to-weight ratios. They are more useful than singular benchmarks, but do not fully encapsulate the full performance of a QPU.

▶ **Application-level benchmarks**

e.g., Q-Score and RACBEM.

These measure the performance of QPUs when solving specific problems. They are similar to classical computing benchmarks such as LINPACK, which is used to rank classical supercomputers. These benchmarks can help compare QPUs vs other QPUs and also compare QPUs with classical computing devices. Application-level benchmarks are analogous to a racecar's lap time at a given racetrack in defined weather conditions. They allow for a limited comparison between competing cars. However, there are different benchmarks that put emphasis on different algorithmic challenges, similar to how a Formula 1 car might be set up to perform well at the Monaco Grand Prix, but would be ineffective on a NASCAR circuit. Manufacturers do not regularly publish application-level benchmarks, as today's QPUs are not capable enough to run sizable applications. As QPUs become more powerful, we expect to be able to track application-level benchmarks in future reports.

Benchmarks

Our QPU dataset was primarily generated through a combination of manufacturer announcements, online searches, and direct queries to QPU providers. The dataset captured a variety of benchmarks. **The key ones are:**

Qubit Counts can be misleading as the number of qubits increases, so too can error rates. This has even prompted some manufacturers to reduce qubit counts to improve overall performance in certain situations. For example, in 2023, IBM followed the release of its 1121-qubit Condor and 433-qubit Osprey QPUs with the higher-performing 133-qubit Heron. Similarly, Quantinuum has been developing and promoting their achievements in higher overall performance on their 20-Qubit H1 QPU², even though their larger 56-Qubit H2 QPU has been commercially available for several years already. This underscores that qubit count alone is not a definitive measure of QPU capability and must be considered alongside other key performance benchmarks.

Coherence refers to how long a qubit maintains its quantum state. Due to interactions with their environment, qubits inevitably lose their quantum information, a process known as decoherence. This is characterized by two timescales: T_1 (energy relaxation) and T_2 (dephasing). T_1 and T_2 are important variables as they dictate the time that calculations can be executed. Trapped ion, and to a slightly lesser extent, neutral atom modalities exhibit T_2 times several orders of magnitude longer than Superconducting qubits, offering an inherent advantage for applications that require longer coherence.

Fidelity: Quantum computers currently experience error rates several orders of magnitude higher than classical systems. These errors arise from imperfections in many areas such as control pulses, inter-qubit couplings, and qubit state measurements, and they reflect the engineering limitations of today's quantum hardware. To characterize and compare quantum error rates, several benchmarks have been introduced. These include single-qubit gate fidelity, two-qubit gate fidelity, readout fidelity, state preparation and measurement (SPAM) error, decoherence-related errors, and crosstalk error. Manufacturers use different terminologies to describe these benchmarks, which can hinder direct comparison (e.g. mid-circuit, median, or average). Each of these benchmarks captures inaccuracies in a specific quantum operation essential to running a quantum circuit. Two-qubit gate fidelity is one of the most critical metrics as it is often the bottleneck in large circuits. It is more prevalent in such circuits and has higher error rates than one-qubit gates.

Quantum Volume (QV) was introduced by IBM in 2017³ and reflects different physical-level benchmarks, such as gate fidelity, qubit count, and connectivity. Unlike the volume of a cube, QV is not computed by simple multiplication, but requires a complex set of statistical tests. QV identifies the largest square-shaped random circuits (where the number of qubits equals the circuit depth) that a

quantum device can implement with high fidelity. QV has been criticized by some scholars and by IBM itself as not being useful for larger devices and also for relying on square circuits which are not typically representative of real-world quantum applications. Despite that, QV has been adopted by many hardware providers and is used in some spec sheets and marketing materials.

Gate speed refers to the time it takes to perform a single-qubit or two-qubit gate operation. It is directly connected to decoherence as together they determine how many operations can be conducted in a system before the qubit becomes

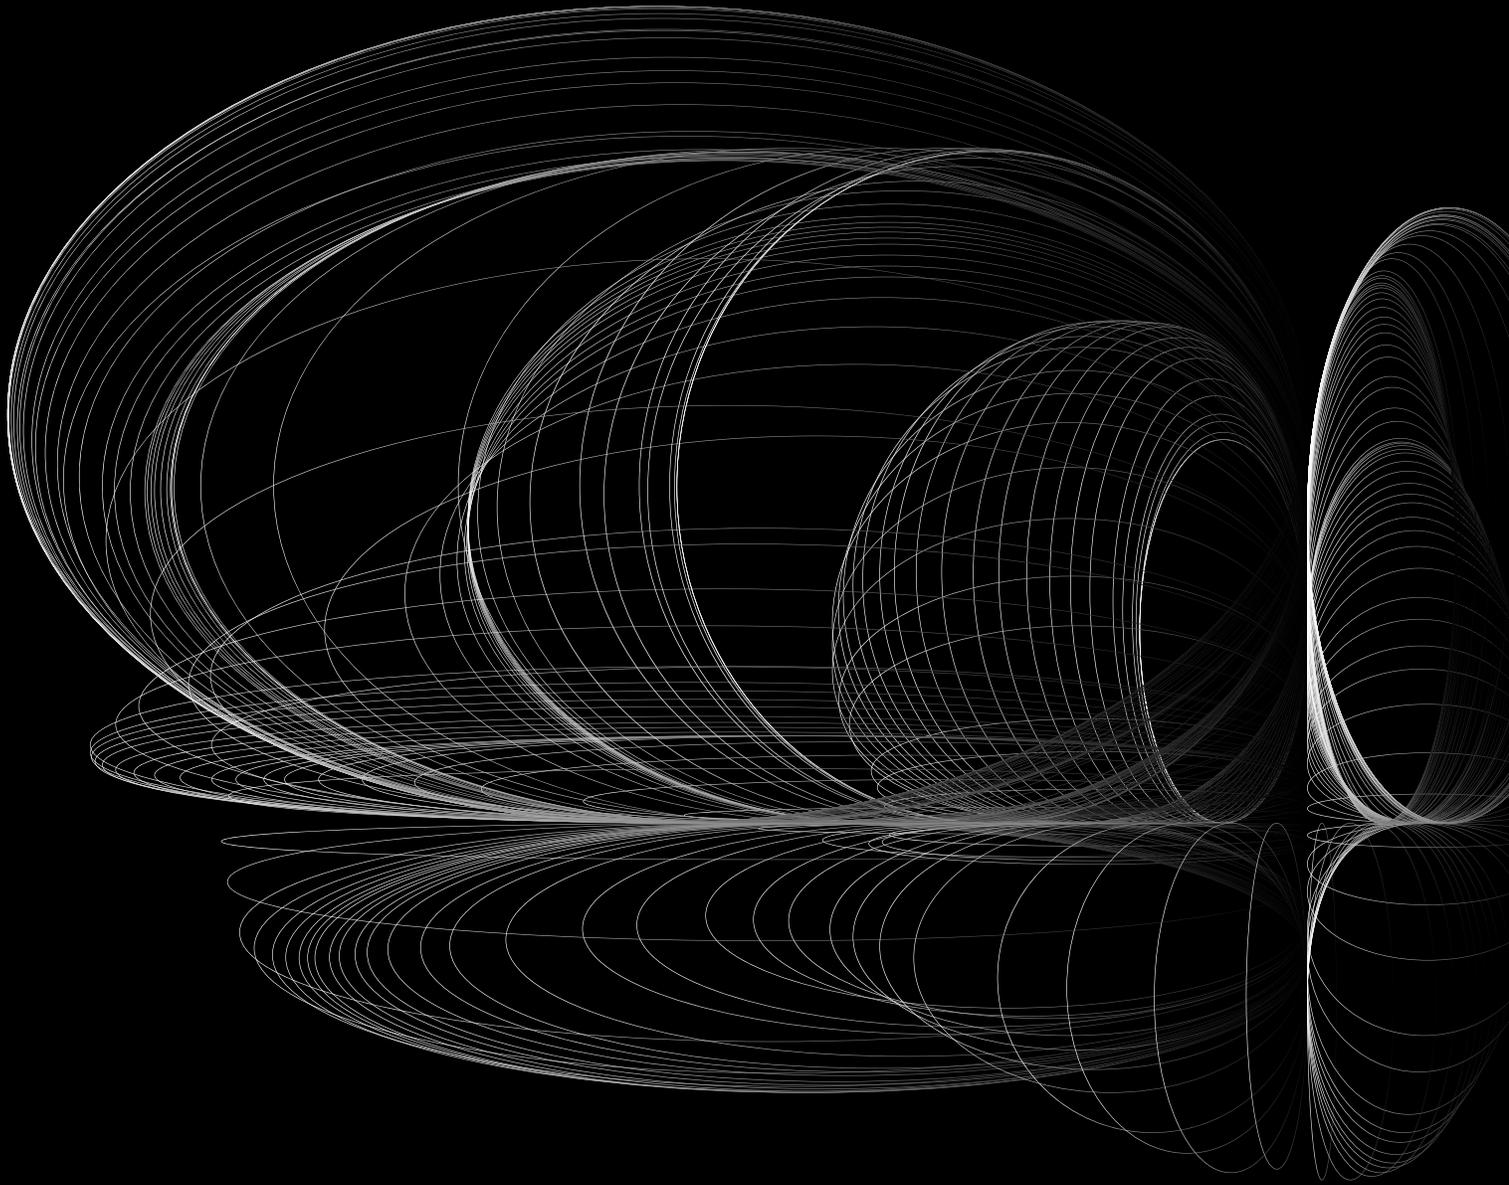
ineffective. Systems with fast gate speeds, such as superconducting qubits and electron spins, often have shorter coherence times than systems like trapped ions or neutral atoms, which have slower gates but much longer coherence times. Current superconducting quantum computers operate at raw gate speeds⁴ in the 1–100 MHz range (and 1–10 kHz when fully burdened with error correction and overhead), these speeds are significantly slower than classical CPUs, which operate at 2–5 GHz. However, such a direct comparison is only partially useful given the fundamentally different computational approaches.

► Execution time

Gate speed is a critical but often underreported metric in quantum computing—many hardware vendors do not disclose it at all. Yet it directly limits the runtime of quantum circuits. For example, consider molecular simulations using a Quantum Phase Estimation (QPE) algorithm, which can require circuits exceeding 10^{13} logical gates. On a Trapped-ion quantum processor, where gate speeds are typically around 10 microseconds, executing such a circuit even once would take⁵ several days. Since a single molecule may require thousands of such full executions to achieve statistical confidence, and since quantum error correction dramatically increases circuit depth and gate count, total runtime could extend into years, well beyond practical limits for most applications. Businesses evaluating quantum computing should estimate execution time based on circuit size, hardware gate speed, and the overhead introduced by error correction. While gate speed imposes a fundamental limit, total runtime can be reduced by optimizing algorithms for parallelism, reducing circuit depth, and improving qubit fidelity to lower the cost of error correction.

► Error correction

Error correction is fundamental to quantum computing. Methods like Surface Codes require an increasing amount of qubits to make physical qubits into a logical one (thousands, or tens of thousands for very large circuits). Google announced an important breakthrough in 2024⁶, demonstrating that their system operates below the fault-tolerance threshold, meaning that adding more qubits and correction cycles leads to a net decrease in logical error rates. This suggests a path forward for a scalable increase of Logical Qubits with a set amount of physical qubits.



Putting a positive spin on it

Composing a detailed QPU list is a challenging task, further complicated by quantum computing vendors often only highlighting their most favorable performance QPU benchmarks. For example, only three out of 31 Trapped Ion QPUs in our dataset reference gate speed in their publicly available specifications. Trapped ions gates are approximately 10,000 times slower than the fastest superconducting gates. IBM is among the most transparent firms when it comes to QPU benchmarking. Most of its relevant performance metrics are publicly accessible, including data for individual QPU instances and even individual qubits.

Modalities

There are several different approaches to designing and operating a specific physical quantum computing system: these are known as modalities. Each modality uses different technological approaches to encoding, manipulating, and reading out quantum information, but result in similar functionality (also called gate-based or measurement-based). Each modality has inherent benefits and weaknesses, reflected in benchmarks such as number of qubits, fidelity, and speed. The table below illustrates the best-in-class commercial or prototyped device from each modality. No clear winner has yet emerged from these modalities.

	Trapped Ions	Superconducting	Neutral Atoms	Photonics	Electron Spin / NV
Qubit Count	56	1121	1125	216	424
2-Q Gate Fidelity	99.91%	99.70%	99.48%	93.80%	99.20%
1-Q Gate Fidelity	99.99%	99.97%	99.90%	99.60%	99.99%
2Q Gate Speed (LogHz)	N/A	7.32	6.18	N/A	N/A
1Q Gate Speed (LogHz)	4.00	7.70	5.39	N/A	N/A
Coherence - T ₁	1,000,00	198	40,000,00	N/A	N/A



A quantum circuit is a sequence of operations (quantum gates) that a QPU follows to solve a problem. It's the foundation of quantum algorithms, which use these circuits to process information.

Superconducting QPUs are electronic circuits created with lithography techniques used for classical computing fabrication. These circuits are cooled to millikelvin temperatures that aid in suppressing thermal noise and allow coherent quantum behavior. They excel in gate speed, qubit count, and have reasonable fidelities, but need to be cooled extensively.

Trapped-ion QPUs implement gate-based quantum computing using individual ions held in place by radiofrequency traps. Gate operations are performed using lasers or microwaves

that manipulate the ions' internal quantum states. Trapped Ions have high fidelity, coherence, and qubit connectivity, but have slower gate speeds and have not yet scaled to large qubit counts.

Photonic QPUs use photons as qubits. Photons propagate through photonic integrated circuits containing linear optical elements such as beam splitters and phase shifters, which can implement certain quantum gates.

Neutral Atom QPUs use atoms—typically alkali or alkaline-earth metals—that are laser-cooled and confined in vacuum chambers using optical and magnetic trapping techniques. While the atoms themselves are ultracold, the hardware operates at near room temperature. Relative to other modalities, they show promise in high qubit counts, but have

slower gate speed and lower fidelity although they are experiencing rapid development.⁷

Electron Spin QPUs leverage the quantum state of single electrons as qubits, offering relatively long coherence times and the potential for high-fidelity control with the added benefit of compatibility with existing semiconductor fabrication techniques. Scalable control

and readout of large arrays remain active areas of research.

Nitrogen-Vacancy (NV) centers in diamond are a promising solid-state platform, where a nitrogen impurity adjacent to a lattice vacancy hosts a localized electron spin used as a qubit. NV center QPUs can operate at room temperature, unlike most other quantum computer hardware.

▶ **Cats, Rails, and Flux**

Superconducting quantum processors are among the most promising modalities. Different sub-architectures are being pursued in designing these QPUs, such as the “Cat qubits” pursued by Amazon and Alice & Bob. The Cat qubit is more difficult to fabricate but increases the fidelity of the qubit as there is only one type of error, phase flip (no bit flip error). Another interesting approach is taken by Quantum Circuits Inc., which is pursuing Dual Rail QPUs that implement a physical redundancy to their Transmon qubits to protect from errors (with a second cavity/rail). Atlantic Quantum is using Fluxonium qubits instead of more traditional Transmon qubits. The Fluxonium qubit shows longer coherence and error correction possibilities—again, at the cost of fabrication complexity.

▶ **Annealers**

A distinct class of quantum computer is the adiabatic quantum computer, also called an annealer, inspired by the metallurgical process with the same name. The principle behind quantum annealing is rooted in the adiabatic theorem, which states that a quantum system will remain in its lowest-energy state if its parameters are changed slowly enough and in the absence of significant noise. Using this phenomenon, an optimization problem can be mapped as an energy landscape of possible solutions with the lowest energy being the best solution. By annealing (i.e., adjusting the system parameters), the system is guided toward the lowest-energy state, which—if reached—yields the optimal solution. D-Wave produced the first commercial annealer in 2010, reaching 128 qubits. Today, the company produces commercial systems with 5,000 qubits. Annealers are treated separately in this report, as their architecture is not directly comparable to gate-based quantum computers. Annealers can achieve much larger qubit counts, but do not implement universal gate-based control. This limits annealers to a narrower class of problems when compared to gate-based QPUs. Only one manufacturer besides D-Wave has announced plans for releasing annealers in the future.

▶ Majorana Qubits

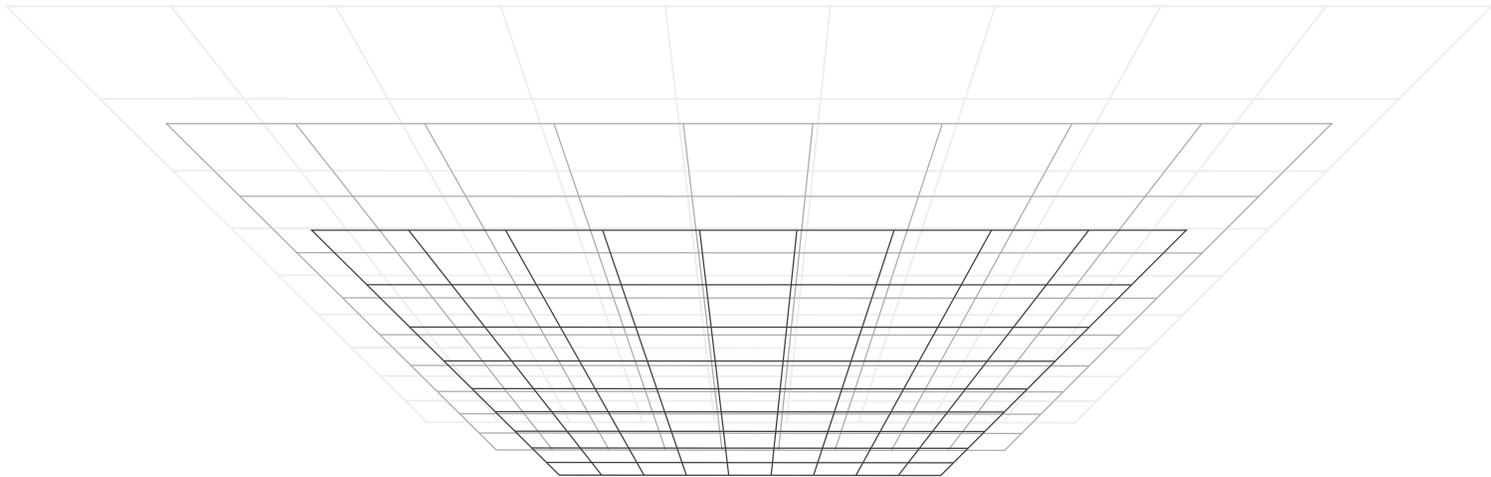
Another type of superconducting qubit is the Majorana qubit, which gained significant attention in early 2025. Microsoft has invested in this approach for over a decade and remains the primary industry player actively pursuing Majorana-based quantum computing. This design uses superconducting nanowires that host Majorana zero modes, exotic quasiparticles predicted to appear at the ends of the wire under specific conditions. Microsoft's Majorana 1 QPU represents a significant milestone toward realizing a Majorana-based quantum processor, though the announcement⁸ (February 2025) has been met with skepticism from parts of the scientific community, as conclusive evidence for the topological nature of the Majorana zero modes remains under debate.

▶ Logical Qubits

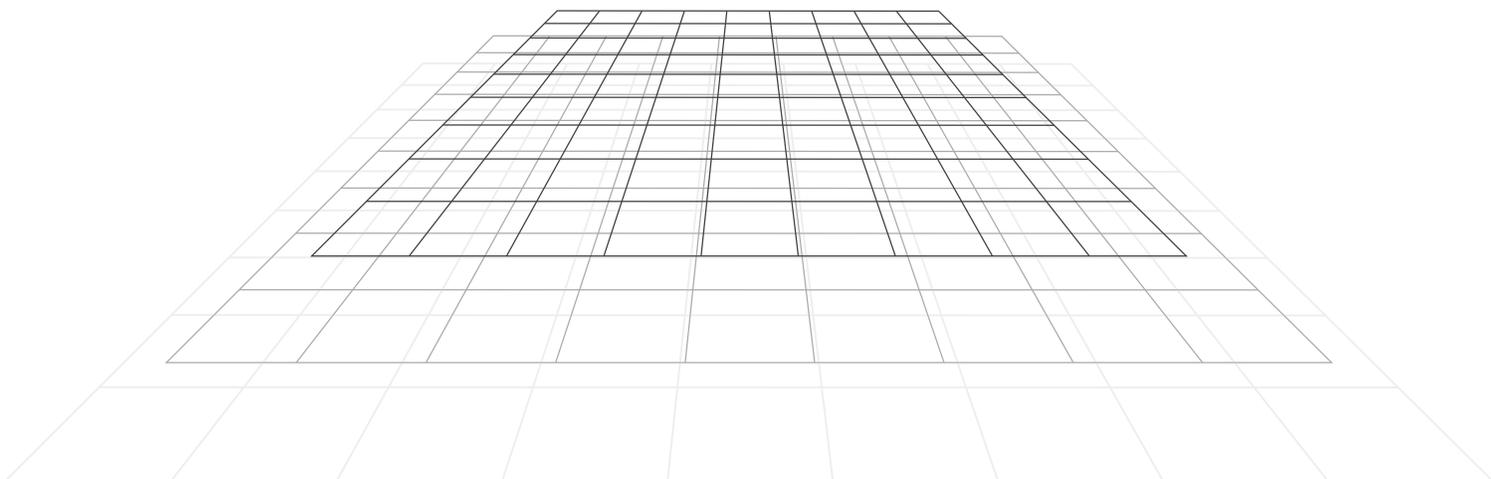
Most quantum algorithms assume that a qubit is “perfect,” i.e., that it behaves perfectly throughout the operations of the algorithm. In reality, qubits are error-prone and short-lived, so we combine many physical qubits using quantum error correction techniques (such as surface codes) to form a more stable unit known as a Logical qubit. Some manufacturers have started using this metric for their QPUs. The term can be misleading, as its practical utility depends heavily on the size and complexity of the circuit it can reliably support. To be viable for applications such as simulating complex molecules, a Logical qubit would need to support circuits millions to billions of gates long—several orders of magnitude beyond current capabilities. As such, when one is presented with a number of Logical qubits for a QPU, the key follow-up question should be: “at what circuit depth?” Only then does the number of Logical qubits convey meaningful information.

▶ Quantum emulators

Since quantum algorithms are inherently probabilistic, they can be emulated by classical computers to a certain level, i.e., run on a classical computer without the need of a QPU. Emulators do not physically utilize quantum effects such as entanglement or superposition. They are particularly useful for testing, debugging, and benchmarking quantum algorithms. Existing classical supercomputers can emulate circuits⁹ with approximately 50 logical qubits. For classical computers, emulating additional qubits becomes exponentially more difficult, while for QPUs, this requires adding incremental logical qubits. Today's best quantum computers are orders of magnitude slower and more expensive to run than the equivalent CPUs.

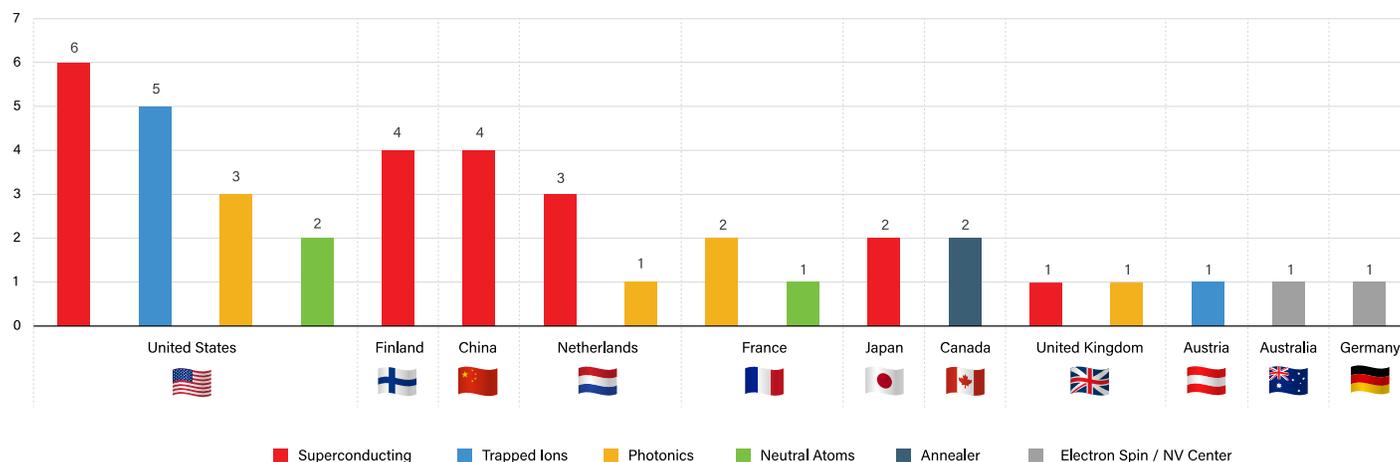


Countries are in a strategic race to achieve high-performance QPUs. The amount of commercially available QPUs globally is in the range of 40 QPUs from two dozen manufacturers. The race is led by the US.



10.1 | QPUs per country and modality

Commercially available QPU models per country



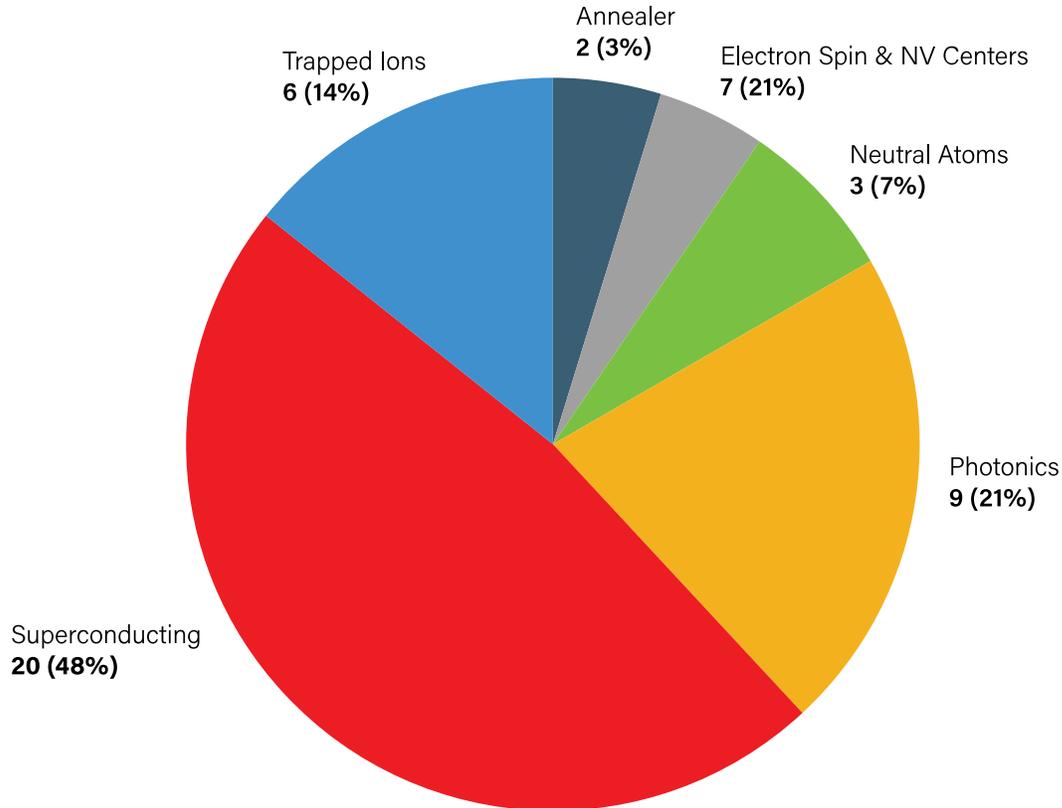
Countries are in a strategic race to achieve high-performance QPUs. The amount of commercially available QPUs globally is in the range of 40 QPUs from two dozen manufacturers. The race is led by the United States, which has the highest number of QPUs and diversity of modalities. China, Finland, and the Netherlands share the second position but their commercial QPUs are of lower performance and smaller in size than the US.

► Data considerations

We list QPUs that are commercially available per country and per quantum computing modality. Country data is allocated based on the location of the manufacturer's headquarters as described in their official materials and website. The QPUs are classified as commercially available if there is public access to the QPU either via on-premise or cloud. This also includes QPUs that may not be available on public clouds, but access is provided to specific partner companies for commercial use, e.g., Google or PsiQuantum QPUs. However, the device must be intended as a useful quantum computer for commercial use and not solely for experimental purposes. The amount of QPUs is determined as uniquely differentiated products actively provided and marketed by the provider, e.g., IBM Eagle and IBM Heron are two distinct QPUs, but Eagle r3-Brussels and Eagle r3-Sherbrooke are considered as one QPU. The amount of QPUs is not necessarily an indication of the progress of each country in quantum computing, as some manufacturers have made several very small QPUs available for basic academic research and teaching, while others have retired smaller but powerful QPUs from their offerings (e.g., IBM).

10.2 | QPUs per modality

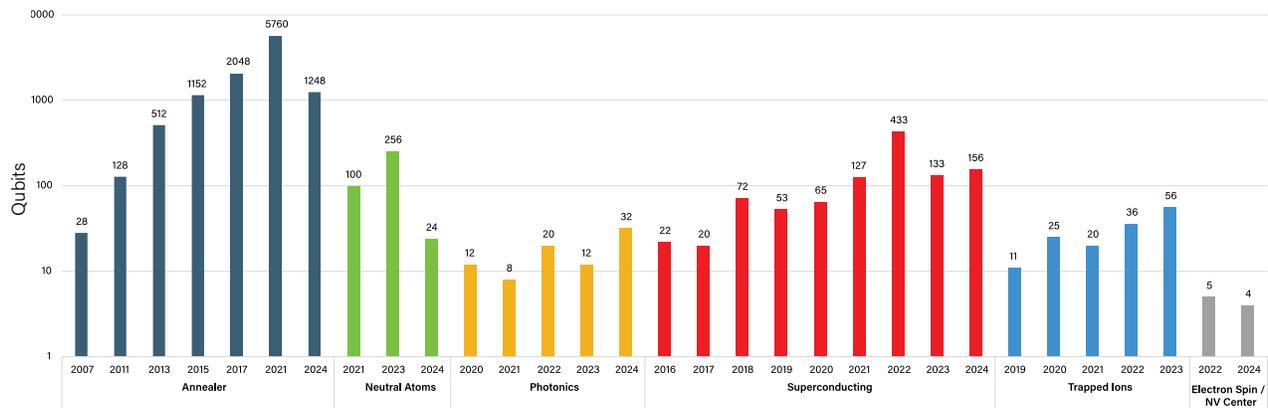
Commercially available QPU models per modality



The leading quantum computing modality is superconducting with more than 40% of commercially available QPUs. This is partially driven by the inherent manufacturing benefits and a historic head start in R&D. However, photonics, trapped ions—and especially neutral atoms and electron spins—are accelerating in quantity and it is expected this trend will continue, while Annealers are becoming increasingly marginalized and NMR QPUs are practically phased out (see Chapter 10.7.1).

10.3.1 | Qubit count per modality

Largest QPU released, prototyped, or planned per year



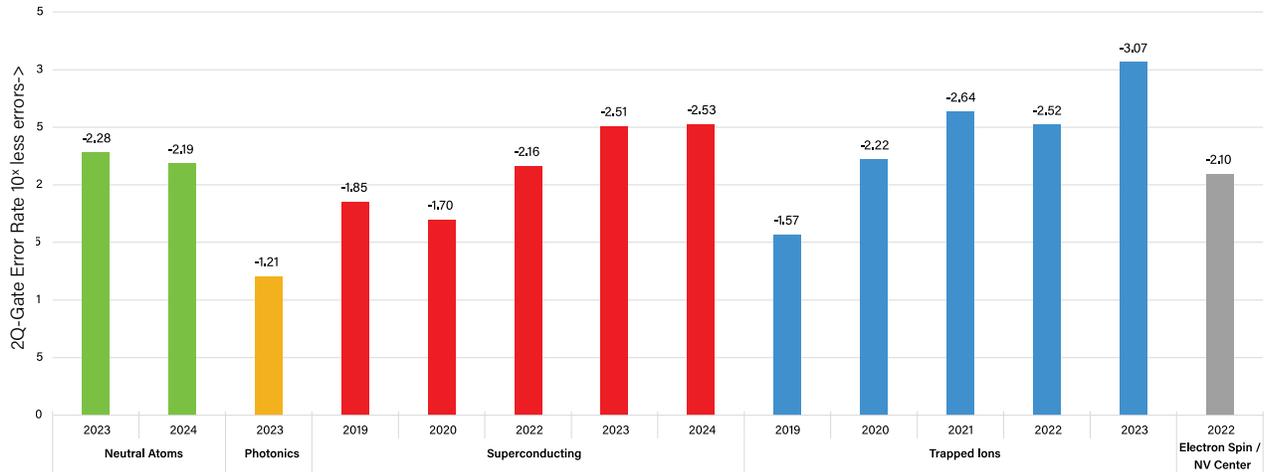
Superconducting QPUs expanded qubit count up to 2022. The more recent decline in absolute numbers reflects an increased focus on improved error correction and higher fidelity (see 10.3.2). Leading Trapped Ion devices are consistently growing their qubit count on an annual basis. The qubit count amongst leading quantum annealers grew steadily across the decade leading up to 2017 but has since stabilized.

► Data considerations

This graph shows the progression of the number of qubits in our dataset over time, considering only the largest QPU announced per modality in that year. For a QPU to be considered, it needed to be officially announced by a manufacturer and made commercially available in the given year (or expanded, e.g., Quantinuum H1 to H1-1). The data does not always show a steady increase, as some calendar years only contained new QPUs that were smaller than previously available.

10.3.2 | Fidelity per modality over time

Best 2Q-gate error rate over time
(available, planned, and prototyped QPUs)



Fidelity for 2-Qubit is a key metric of performance improvement. Trapped Ions have shown consistent growth and demonstrated the highest overall fidelity. Superconducting QPUs experienced a decline in top fidelity from 2018 to 2022 until peak fidelity was achieved by the Alibaba QPU. Alibaba subsequently withdrew from the quantum computing market, and IBM and Google caught up to similar fidelity rates in 2023 and 2024. Photonics and NV Center QPUs are still relatively nascent and haven't achieved the top-performing fidelities of trapped ion and superconducting QPUs.

► Data considerations

This graph shows the progression of fidelity levels in new, commercially available QPUs in our dataset over time. We only considered the highest fidelity announced per modality per year. Error rates are given in a log10 scale, i.e. -3 translates to a 0.001 error rate which corresponds to a 99.9% or 0.999 fidelity. The noted error rates should be treated with caution as there are significant differences in the way they are measured for each QPU, e.g., mid-circuit vs first-gate-measurement, average vs median of several measurements across qubits, different gates (CZ, SWAP, etc.), and different iterations of the same QPU that give different values. In the case of conflicting values, we followed the data mismatch process detailed in the methodology chapter.

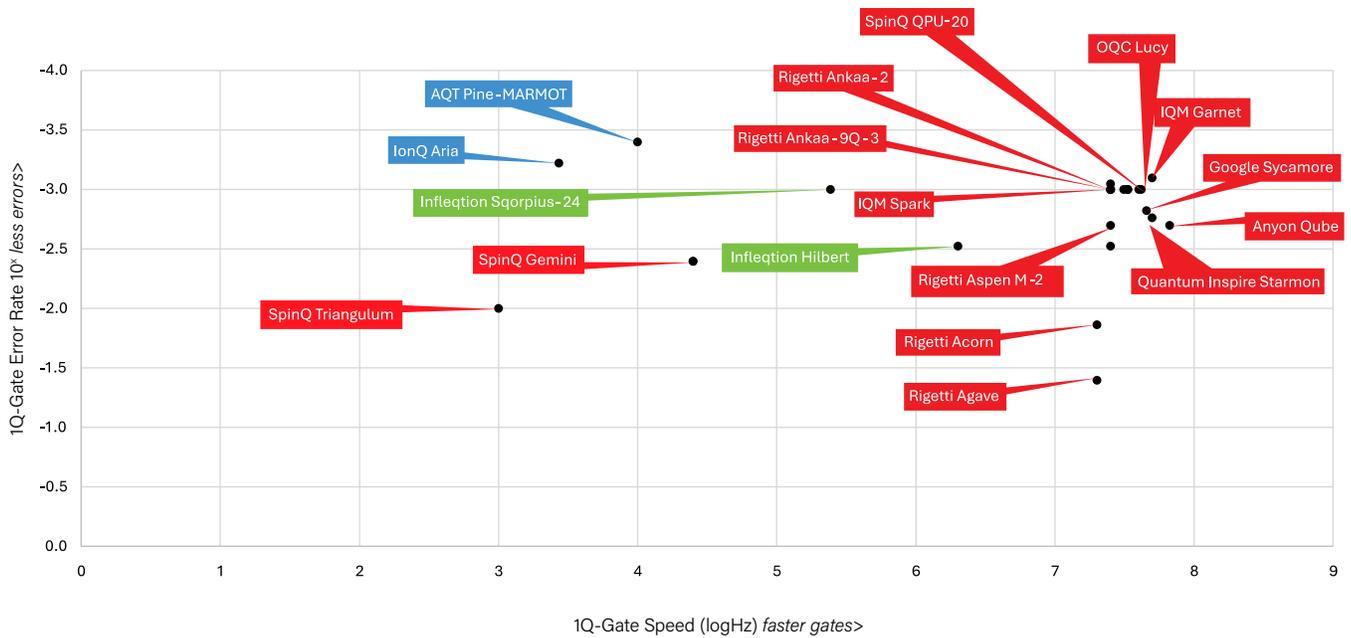
10.5 | Gate time versus gate fidelity



To determine the maximum length of a circuit for a QPU, an important metric is the comparison of the speed of executing a single gate to how accurate this gate is (fidelity). For real-life scenarios such as Shor's Algorithm for decrypting information utilizing RSA2048, more than 10^{13} Logical Gates are required. Slow gate speeds at that size lead to calculation times of days or even months with some modalities.

1-Qubit and 2-Qubit Gates devices are identified above; the latter is more interesting as they are more common in large circuits for most algorithms. The superconducting IBM Heron and IQM Garnet are the class leaders. Notably absent are ion traps and neutral atoms QPUs, as manufacturers tend not to disclose exact gate speeds, which are expected to be orders of magnitude slower than superconducting QPUs (as can be seen in the 1-Qubit graph).

1-Qubit gate time vs error rate

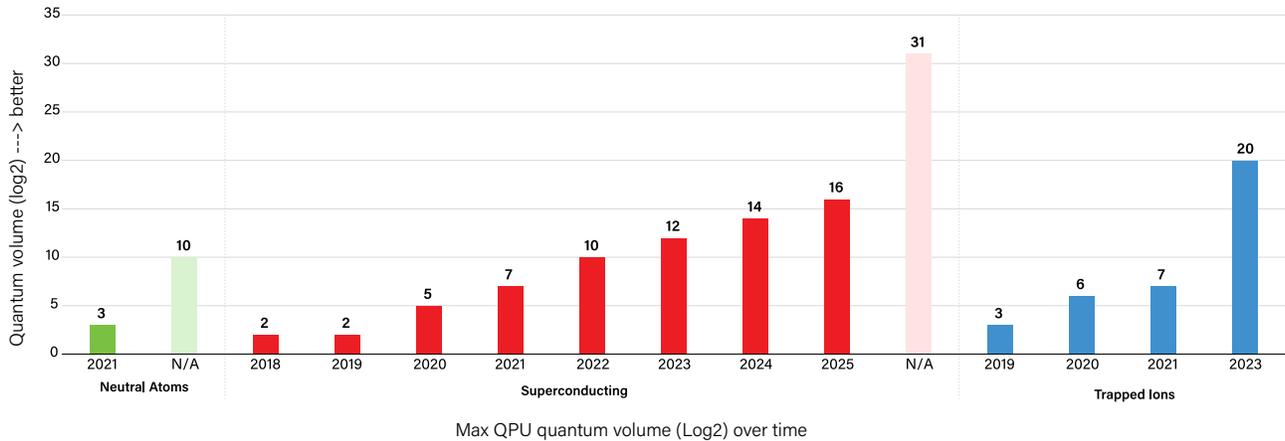


► Data considerations

We chose 2-Qubit-Gates like CZ, CNOT, and SWAP as they are used for most quantum algorithms and make up the majority of gates for these circuits. The 2Q Gate time is the time required for the execution of a 2-Qubit gate and is given in Hertz (Hz) in a logarithmic scale, where higher values mean faster gate speeds, i.e., 7 logHz corresponds to a 100ns gate speed and 3 logHz to 1,000,000ns. The error rates are given in a log10 scale where higher is better, i.e., 10⁻³ translates to a 0.001 error rate, which is 99.9% fidelity. The datapoints with missing labels in the graphs are closely related QPUs (e.g., different instances of IBM Eagle). To illustrate the performance comparison, we included the 1-Qubit Gate graph, which shows that these would likely land on the top left quadrant, trading high fidelity for low gate speeds.

10.6 | Quantum Volume

Max QPU Quantum Volume (Log2) over time



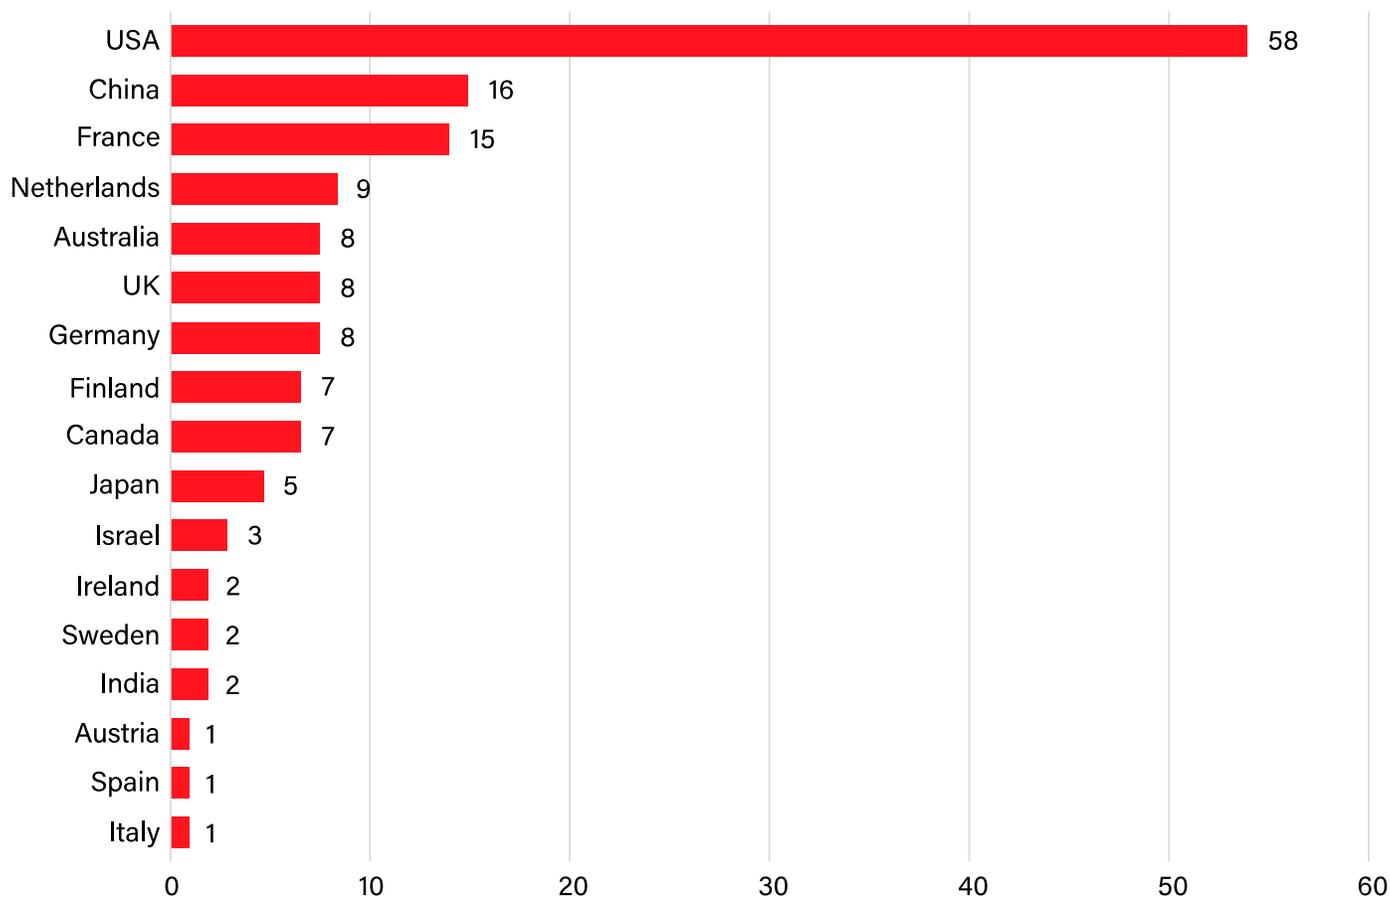
As described in an earlier section, Quantum Volume (QV) is an aggregated metric designed to reflect a more holistic view of overall performance of a QPU. Although its usefulness is heavily debated, manufacturers have been claiming this metric. There is a substantial and consistent increase in QV performance of the top-performing QPU across modalities. However, the stated QV values are hard to validate, and in some cases the only data has been manufacturer claims.

► Data considerations

The Quantum Volume (in log2 basis) is listed for each QPU at the given year. QV is itself a debated benchmark (even by IBM themselves), which has been used less in the last few years. Despite all these caveats, we chose to include QV, as out of all aggregated metrics (like RACBEM, Algorithmic Qubit, CLOPS, etc.) it is the one that has published values for a sizable amount of QPUs, and as such a progression on quantum computing capabilities can be roughly traced over time. To note is also that the values under 2025 and "N/A" are manufacturer plans, not yet available QPUs.

10.7.1 | A look into the future: QPUs per country and modality

Planned, prototyped, and available QPUs per country

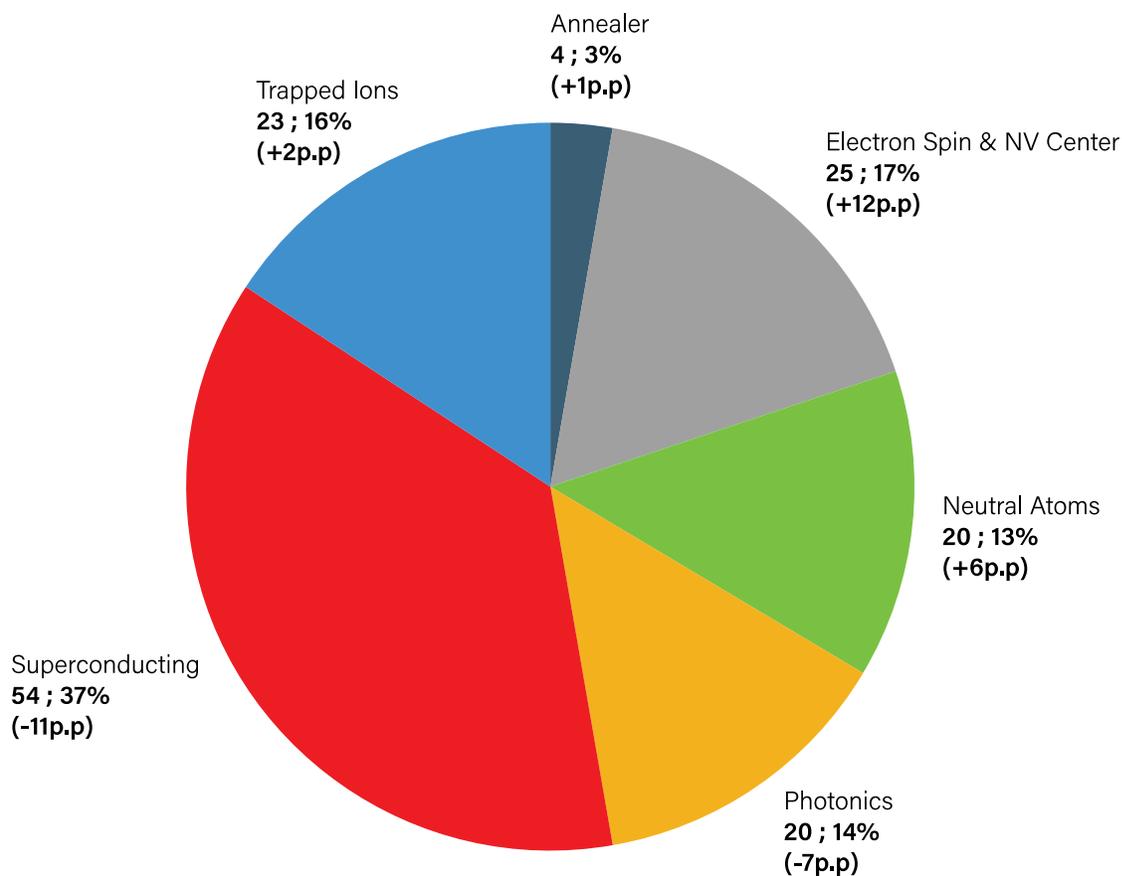


Quantum computing remains in the early stages of technical development. Numerous startups and established companies are actively developing prototypes and announcing roadmaps for future QPUs. A comprehensive overview of these efforts provides valuable insight into the evolving technological landscape and strategic directions within the field. The US is the clear leader in the number and diversity of QPUs announced. China is in second place, but is closely followed by France. The Netherlands, Germany, Australia, Canada, Finland, and the UK have announced 7-10 QPUs each. This data gives an approximate measure of country activity but is not an indication of QPU quality. It is also an artifact of QPU planned announcements and does not fully reflect the probability of their successful release.

10.7.1 | A look into the future: QPUs per country and modality

QPUs commercially available, prototype, or planned

per modality; change in percentage points to current distribution of QPUs



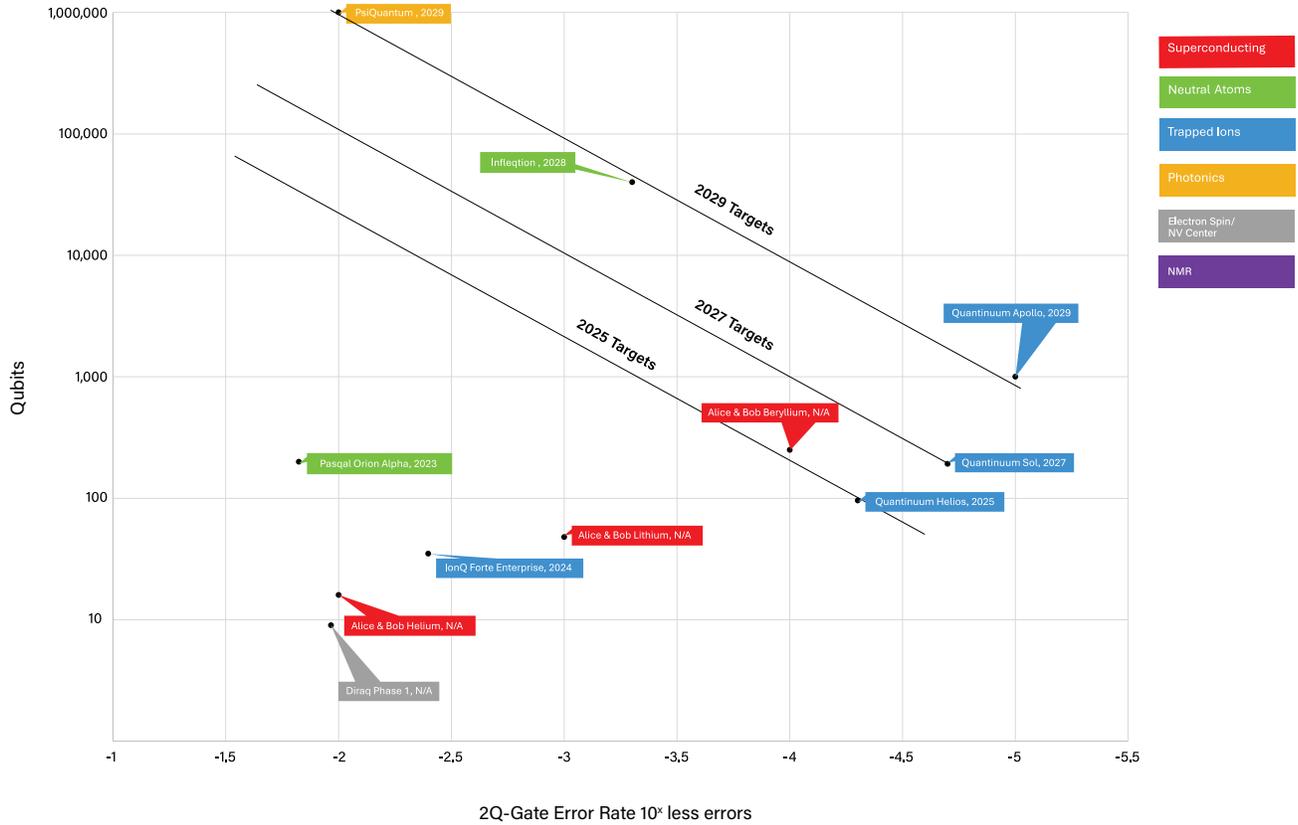
Looking at the future for different modalities, electron spin, NV Centers and neutral atoms are planned to become increasingly prevalent while NMRs and annealers are stagnant and may be phased out. Photonics, superconducting, and trapped ion QPUs may have lower overall shares in the future due to the higher growth levels of other modalities.

► Data considerations

The data in these graphs includes prototype devices, which are not intended for commercial usage and are not available to the wider community of researchers but are used by a manufacturer for research to develop a new product. It also includes Future Planned QPUs, which are announced in a manufacturer roadmap or interview. Due to the dynamic evolution of startups in the space, recent announcements and changes in QPU roadmaps may not be fully captured in our dataset. The amount of QPUs is not necessarily an indication of the progress of each country in quantum computing, as some manufacturers have made several very small QPUs available for basic academic research and teaching, while others have retired smaller but powerful QPUs from their offerings (e.g., IBM).

10.7.2 | A look into the future: qubit count and fidelity

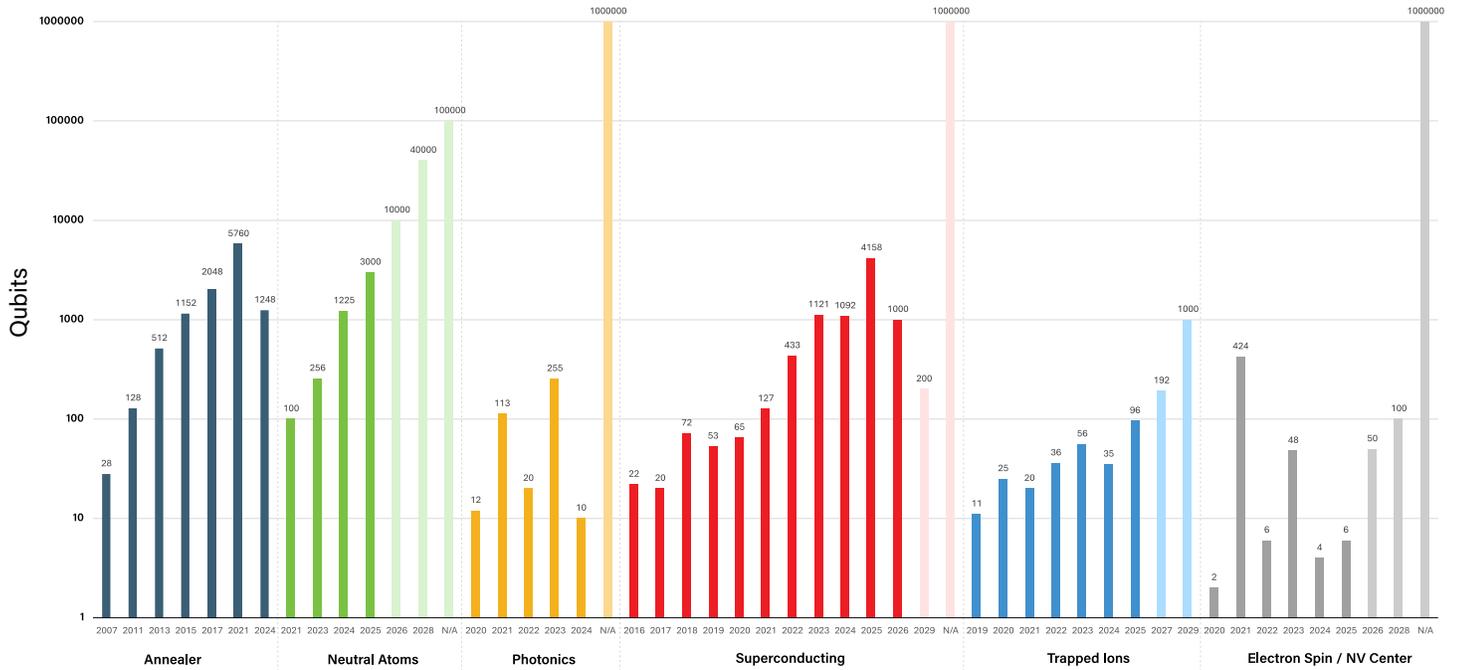
QPUs vs 2Q gate fidelity Announced QPUs



QPU manufacturers are beginning to share more forward-looking roadmaps, offering insights into how their systems might compare across benchmarks. PsiQuantum leads in projected qubit count, Quantinuum excels in error rates, and Infleqtion positions itself as a balanced performer across both metrics.

10.7.2 | A look into the future: qubit count and fidelity

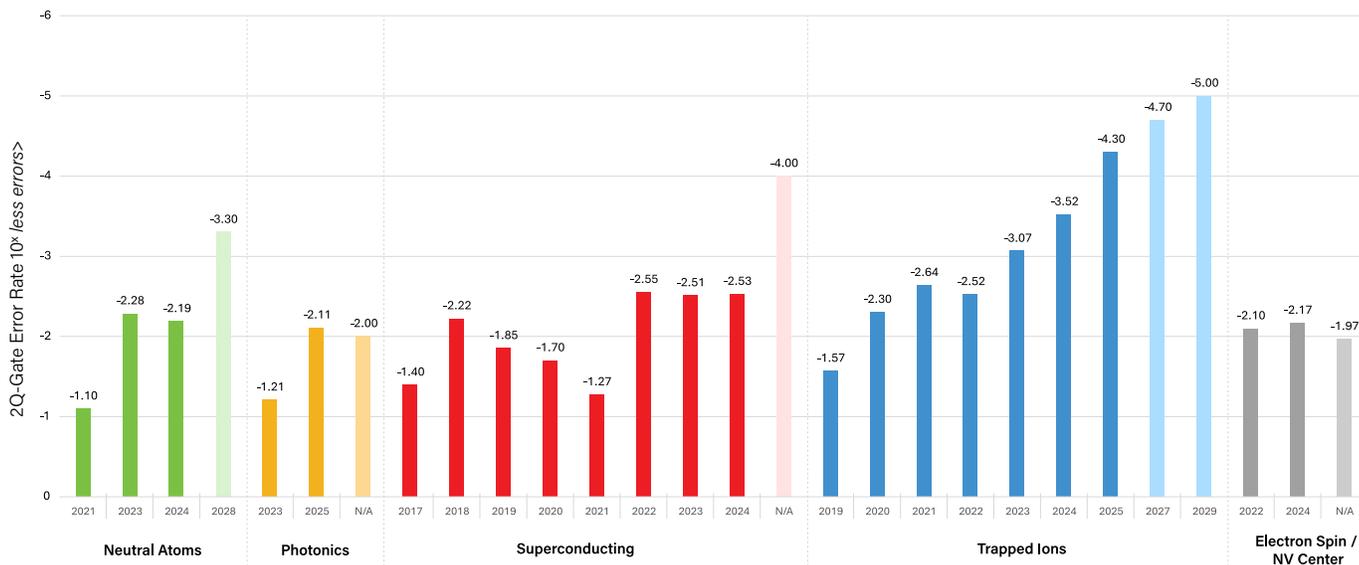
Largest QPU released, prototyped, or planned per year



While qubit counts are expected to continue rising, the pace of growth is moderating. This reflects a shift in focus toward improving performance through better error correction and higher qubit fidelity rather than simply scaling up qubit numbers.

Best 2Q-gate error rate over time

Available, planned, and prototyped QPUs



Trapped ion systems aim for exponential gains in fidelity and are on track to continue outperforming other modalities in that benchmark. Neutral atom platforms also show strong ambitions in this area, while other technologies appear more conservative in their likely fidelity trajectories.

► Data considerations

These graphs show the progression of the fidelity and qubits counts for published QPUs over time, including any future plans, considering only the largest QPU announced per modality in that year. The data does not demonstrate a constantly increasing trend, as some calendar years saw smaller QPUs than previously available. Although approximately 60 manufacturers have announced approximately 90 future QPU models, only 11 QPUs have been provided with both target qubit count and fidelities, which is why there are fewer QPUs in the first graph, QPU vs 2-Qubit Fidelity.

10.8 | Future research

Benchmarking is an important exercise in the advancement of our understanding of quantum computing technology as it enables informed decision-making and supports the longer-term goal of standardized comparisons. Our contribution of a publicly accessible overview aims to improve general transparency and allows researchers and community members to engage in a more detailed dialogue regarding the performance of various systems. We encourage industry members and other stakeholders to contribute to these goals by adding their data on an ongoing basis. This will help bridge the gap in this domain, where standardized datasets have been scarce. This is a rapidly and constantly evolving space. By keeping this resource updated and relevant, we are hoping to foster further collaboration and innovation.

You can reach us at contact@qir.mit.edu.

► Footnotes

¹ The difference of 200 indexed QPUs vs 40 commercially available in approximate numbers: 40 are retired, 30 are prototypes but not commercially accessible, and 90 are planned i.e., not released. To be commercially available, they have to be accessible via cloud or on-premise (10 out of 40 QPUs in our dataset are on-premise).

² 'Quantinuum Extends Its Significant Lead in Quantum Computing, Achieving Historic Milestones for Hardware Fidelity and Quantum Volume' <<https://www.quantinuum.com/blog/quantinuum-extends-its-significant-lead-in-quantum-computing-achieving-historic-milestones-for-hardware-fidelity-and-quantum-volume>> accessed 3 April 2025.

³ Bishop, L. S., Bravyi, S., Cross, A., Gambetta, J. M., & Smolin, J. (2017). Quantum Volume.

⁴ Olivier Ezratty, 'Understanding Quantum Technologies 2024' (Opinions Libres - Le blog d'Olivier Ezratty) <<https://www.oezratty.net/wordpress/2024/understanding-quantum-technologies-2024/>> accessed 3 April 2025.

⁵ Raffaele Santagati and others, 'Drug Design on Quantum Computers' (2024) 20 Nature Physics 549.

⁶ Rajeev Acharya and others, 'Quantum Error Correction below the Surface Code Threshold' (2025) 638 Nature 920.

⁷ Simon J Evered and others, 'High-Fidelity Parallel Entangling Gates on a Neutral-Atom Quantum Computer' (2023) 622 Nature 268.

⁸ 'Microsoft's Majorana 1 Chip Carves New Path for Quantum Computing' (Source) <<https://news.microsoft.com/source/features/innovation/microsofts-majorana-1-chip-carves-new-path-for-quantum-computing/>> accessed 3 April 2025.

⁹ Thomas Häner and Damian S Steiger, '0.5 Petabyte Simulation of a 45-Qubit Quantum Circuit', Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (2017) <<http://arxiv.org/abs/1704.01127>> accessed 3 April 2025.

11 | Appendix

Chapter 9 | Quantum networking testbeds

The dataset was created by merging input from the Center for Quantum Networks (CQN) researchers, QIR professional network and publicly available information by GQI Quantum Computing Report (accessed in June 2024).

Chapter 10 | Quantum processor benchmarking

The dataset of QPUs was composed by a combination of a keyword-based online search and official announcements, references to QPU lists made available to us, and direct query to QPU manufacturers. The data was collected from January 2024 to April 2025.

In particular, a list of known manufacturers was created based on the sources of The Quantum Insider, Olivier Ezratty, and Wikipedia. For each manufacturer, the official website was interrogated to retrieve the indicated benchmarks. For datasets not on manufacturer's websites, we utilized web searches (Google) for official announcements from manufacturers and related news articles.

Additionally, scholarly articles were identified via Arxiv and Google Scholar using the following keywords for benchmarks: Quantum Volume, CLOPS, EPLG, Q-Score, benchmarking.

During this process, additional manufacturers/QPUs were identified and added to the QPU list. Lastly, each manufacturer was contacted for verification of records—either to an existing contact of the QIR team, or to the communications address listed on manufacturer's website. The final list was reviewed by the QIR team and experts in their professional network.

► Footnotes

¹ Gaida, J., Wong-Leung, J., & Robin, S. (2023). Critical technology tracker. Who Is Leading the Critical Technology Race. A Project by the Australian Strategic Policy Institute. <https://techtracker.aspi.org.au>



Quantum Index Report 2025
MIT Initiative on the Digital Economy

245 First Street, Room E94-1521, Cambridge, MA 02142 USA

qir.mit.edu | ide.mit.edu