

Document Version

Final published version

Licence

Dutch Copyright Act (Article 25fa)

Citation (APA)

Awschalom, D. D., Bernien, H., Hanson, R., Oliver, W. D., & Vučković, J. (2025). Challenges and opportunities for quantum information hardware. *Science (New York, N.Y.)*, 390(6777), 1004-1010.
<https://doi.org/10.1126/science.adz8659>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

A QUANTUM WORLD

REVIEW

Challenges and opportunities for quantum information hardware

David D. Awschalom^{1,2*}, Hannes Bernien^{1,3,4}, Ronald Hanson^{5,6}, William D. Oliver^{7,8,9}, Jelena Vučković^{10,11}

Quantum technologies have made impressive progress over the past decade. In some areas, such as quantum sensing and key distribution, these technologies are moving from the laboratory to enable real-world applications. However, for areas such as quantum computing, entanglement-enhanced sensing, and a global quantum internet, we are in an equivalent of the early transistor age, and hardware breakthroughs are required in multiple arenas to reach the performance necessary for the envisioned applications. In this Review, we assess the current state of the art of quantum information hardware and identify key challenges and opportunities ahead. We draw inspiration from the history of scaling and development of classical electronics and photonics to anticipate progress in the field.

Recent progress in the fields of quantum information science and engineering includes both rapid advances in fundamental research and practical demonstrations of quantum computers, networks, and sensors. As was the case with microelectronics and photonics, the strong partnerships between academia, industry, and government have been instrumental in fueling many of the key breakthroughs that ultimately enabled scaling, adoption, and commercialization. Quantum technologies will similarly face many such challenges, some of which are familiar to classical technologies (e.g., the tyranny of wires) and others that are specific to quantum technologies (e.g., sustained, quantum-coherent control of large systems). Moreover, it is likely that advanced instrumentation for fabrication and characterization will need to be developed to reach these goals. These challenges must be met to deliver meaningful quantum advantage at scale for current and unforeseen applications.

What lessons could be learned and what expectations set for quantum information technologies from our experiences with classical (von Neumann model) computing? The cost per compute operation in MIPS/\$ (millions of instructions per second per dollar) has dropped more than 14 orders of magnitude over the past 75 years (*1*), and the number of transistors on a processor chip has increased from a few thousand in the early 1970s to almost 100 billion in systems on a chip today. On the hardware side, this advancement was accomplished by radical changes in base technology, from vacuum tubes to silicon transistors (1950s), to integrated circuits (1960s to '70s), and the move from bipolar transistors to complementary metal-oxide semiconductor (CMOS) transistors (1990s). In each case, progress was driven by innovations in materials, devices, and circuit design. A few broad principles that guided this success may be valuable for consideration in quantum technology research.

First, once the underlying physics and device concepts were established, the progress had a notably top-down-driven methodology, meaning that system-level considerations defined the specific circuit, device, and materials requirements for forthcoming innovations. The ensuing discovery-based science research thus had clear guidelines that scientists could target and map onto emerging technology. The replacement of the traditional silicon dioxide gate insulator in the transistor structure with new hafnium dioxide-based dielectrics in the early 2000s is an example, and 3D NAND Flash is another. These efforts were successful because of coordination of expertise and code-sign across the technology stack. On the other hand, bottom-up approaches, although essential in the early stages of a technology for discoveries and the understanding of the fundamental principles, have proven less successful as a design strategy. For instance, materials discoveries that then sought an application for novel properties were not as successful, because of a lack of codesign perspectives.

Second, the community (industry, academia, and national facilities) spent considerable effort in developing a shared and fundamental understanding of the underlying science that was disseminated by refereed publications, which then guided future research and development. This collective body of knowledge aided the entire research community. For instance, a firm understanding of the physics of scaling and excitation transport in transistors led to Dennard's scaling law (of which Moore's law is a consequence) and, more recently, the move to nonplanar transistors such as FINFETs (fin field-effect transistors) and nanosheet FETs. Much of this science was considered pre-competitive and avoided the inefficiencies of R&D that is prematurely siloed as confidential.

Third, materials and process science innovations underpinned almost all major progress throughout the history of classical computing. Pfann's semiconductor purification technique (1950s) led to practical transistors, ion implantation permitted Moore's law scaling, extreme ultraviolet (EUV) lithography enabled extreme scaling of transistors, and the advent of atomic layer deposition (ALD) made nonplanar transistors possible. Finally, patience has been a key element in many of these landmark developments and points to the importance of tempering timeline expectations in quantum technologies. For example, EUV materials and optical systems work began in earnest through a partnership between the US Department of Energy labs and industry in the mid-1990s (*2*), but it was only in 2019 that EUV, essential for building today's high-performance transistors, was first used in manufacturing. Similarly, ALD, one of the most important thin film deposition methods used today because of its ability to conformally coat nonplanar surfaces, was developed in the 1960s and 1970s, but it was only in the 2000s that it saw use in microelectronics manufacturing.

Although it is not a given that quantum technologies will develop similarly to classical computing, important lessons can be learned. We consider the progress and opportunities in emerging hardware for quantum information processing. How do we progress from current state of the art to scalable technologies? And what are the challenges that the community may need to address and how might we address them to reach this goal?

Current state of the art

Quantum technologies today are transitioning from laboratory curiosities to technical reality. Assessing progress of these new technologies is challenging, as their future main application domains and their

¹Pritzker School of Molecular Engineering, University of Chicago, Chicago, IL, USA. ²Department of Physics, University of Chicago, Chicago, IL, USA. ³Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Innsbruck, Austria. ⁴Institute for Experimental Physics, University of Innsbruck, Innsbruck, Austria. ⁵QuTech, Delft University of Technology, Delft, Netherlands. ⁶Kavli Institute of Nanoscience, Delft University of Technology, Delft, Netherlands. ⁷Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁸Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁹Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA, USA. ¹⁰Ginzton Laboratory, Stanford University, Stanford, CA, USA. ¹¹Department of Electrical Engineering, Stanford University, Stanford, CA, USA. *Corresponding author. Email: awsch@uchicago.edu

respective figures of merit are not yet fully known. Identifying and agreeing upon such metrics is an important challenge for the field in the coming years, benefiting researchers, developers, public and private investors, and the broader community. Lacking such established metrics today, we present in Fig. 1 the six most advanced quantum technology platforms and use the framework of “technology readiness levels” (TRLs) for quantum computing, quantum simulation, quantum networking, and quantum sensing as a proxy for technological maturity. To provide an external, aggregated view on the field, TRL values were taken directly from queries on ChatGPT and Gemini.

We would like to emphasize two important points for interpreting the TRLs. First, the original TRL mapping requires defining concepts such as prototype and operational environment, which will depend on the not yet fully known application domain. The generated output therefore mainly provides a snapshot for relative maturity between the different platforms rather than absolute progress toward the most advanced applications.

Second, achieving a higher TRL indicates that a higher level of system sophistication has been demonstrated, even if the raw performance of the platform is still consistent with an early-stage technology. For example, the Intel 4004 processor (shown in Fig. 2) was at a rather high TRL for 1971 and indeed led to commercial applications, but its performance was minimal compared with the CPUs (central processing units) and GPUs (graphical processing units) available to us today, 50 years later. In this sense, a high TRL for quantum technologies today does not indicate that the end goal is achieved, nor does it indicate that the science is done and only engineering remains. Rather, a high TRL for an emergent technology reflects a system-level demonstration in an operational environment that can (and must) be improved going forward. For instance, although today’s cloud-accessible quantum computer prototypes are at relatively high TRL for the use case of research and education, they are still far away from the error-correcting machines that are required for commercial quantum utility.

Likewise, increasing TRL takes time. The first silicon ring modulators were demonstrated 20 years ago with modest TRL (3), and their performance and TRL have since improved substantially. Today, ring modulators are densely integrated with other photonics and electronics on circuits fabricated in commercial foundries, as shown in Fig. 3.

Similarly, quantum technologies in 2025 are advancing rapidly, and sophisticated systems demonstrations are being conducted, even in operational environments. However, as emergent early-stage technologies, their raw performance remains relatively nascent and must be considerably improved to reach their full promise. For example, implementing factoring (4, 5) or quantum chemistry algorithms (6) at scale is estimated to require thousands to tens of thousands of




Fig. 1. Snapshot of technology readiness levels of emergent quantum platforms. Comparison of TRLs of emergent quantum platforms across key application areas as assessed by ChatGPT and Gemini (see supplementary materials). TRLs range from 1 (basic principles observed in a lab environment) to 9 (emergent platform proven in an operational environment). A TRL 7–9 platform may not yet be highly performant, but it has been transitioned out of the lab into a system-level, operational environment, even though the system itself may today be relatively immature compared with a desired future outcome. The range of TRLs output by ChatGPT and Gemini reflects, in part, the variation in subfields within each category.




Fig. 2. Historical perspective of scaling classical electronic processors. (Left) Intel 4004, one of the first commercial integrated circuits, with 2300 transistors, 4-bit instruction set targeting integer arithmetic (not floating point), 10- μ m lithography node, and 750-kHz clock capable of around 5×10^4 to 10×10^4 instructions per second. [Photo of Intel C4004 by Thomas Nguyen (https://commons.wikimedia.org/wiki/File:Intel_C4004.jpg), CC BY-SA 4.0] (Right) NVIDIA Blackwell GPU with a total of 208 billion transistors (two dies, each with 104 billion transistors), up to 32-bit instruction set targeting both integer and floating point operations, 4-nm lithography node, and a 2- to 3-GHz clock capable of 14 to 18 petaFLOPS (floating point operations per second) when executing 4-bit floating point (FP4) AI inference. [Source: NVIDIA]

fault-tolerant logical qubits (i.e., millions of physical qubits operating well beyond 99% gate fidelity) executing billions to trillions of gate operations. Realizing such mature and sustained computation with manageable overhead will require improved performance throughout the entire system—not just in the gate fidelities—and is very far beyond what is achievable with today’s early-stage quantum computers (see Fig. 2 for a classical computing analogy). Similarly, whereas quantum communication allows for a few commercial applications, such as quantum key distribution, with current photonics technology, unlocking the full potential of new quantum-enabled privacy and




Fig. 3. Historical perspective of scaling silicon photonics. (Left) Early silicon photonics: gigahertz-speed ring resonator modulator. [Adapted with permission from Michal Lipson, Columbia University; (3)] (Right) State-of-the-art photonics, showing high-density integration of ring modulators (zoom ins). The circuit is a 10 Tb/s optical interconnects chip fabricated in GlobalFoundries 45CLO process. [Figure courtesy of Firooz Aflatouni, University of Pennsylvania]

communication functionalities (7) will require quantum networks linking qubit registers of size and performance well beyond current demonstrations (8). Next, we briefly summarize the current state of the art for several technology platforms.

Superconducting qubits and quantum dots

Superconducting qubits (9) and lithographically defined quantum dots (10) are both examples of “artificial atoms”—manufactured electrical structures that behave quantum mechanically at dilution refrigerator temperatures. Superconducting qubits are anharmonic electrical oscillators built from inductors, capacitors, and Josephson junctions that feature discretized states of flux, charge, or “phase.” Semiconductor quantum dots use lithographically defined gates with applied voltages to trap single electrons, with the electron charge and spin states serving as the qubit. Silicon fabrication technology is utilized for the manufacture of devices for both types of qubits, albeit with careful processing windows and materials to realize coherent quantum devices.

State-of-the-art superconducting qubit processors have >100 qubits arranged in a two-dimensional (2D) array, with some processors also using a similar number of tunable couplers to turn on/off interactions between neighboring qubits. In such large arrays, the gate fidelities are in the approximately 99.95 to 99.99% range for single-qubit operations and in the approximately 99.5 to 99.9% range for two-qubit operations (11, 12). Isolated devices can do about a factor of 5 to 10 better (13, 14).

One advantage to superconducting qubits is their speed: Qubit operations are performed in about 10 to 40 ns (15), and readout can be performed with >99% fidelity in 100 to 200 ns (16, 17). This speed leads to quantum error correction cycle times of around 1 μs (15, 18). In a recent demonstration, increasing the error correction code distance from 3 to 5 to 7 (corresponding to surface-code logical qubits with 17, 49, and 97 physical qubits, respectively) reduced the error rates by more than a factor of two with each step, meaning: more qubits, lower error. At distances 5 and 7, the logical error rate was lower than even the best physical qubit (12). In the present “brute force” scaling era, where there are one to three wires per qubit, approximately 1000 to 10,000 qubits and their associated wiring can fit inside commonly available dilution refrigerators (19). To scale to larger numbers of qubits, such refrigerators would need to be networked together (20) and/or larger ones would need to be manufactured.

Spin qubits in semiconductors can be hosted in both gate-defined quantum dots, such as in Si/SiGe, Ge/SiGe, or Si-metal-oxide-semiconductor devices, and by atomically patterned dopants (optically interfaced spin qubits are discussed separately in the next subsection). Their size is set by lithographically patterned electrodes that can be fabricated by industrial processes and drive quantum gates at rates approaching gigahertz. The small size is a strong feature in terms of footprint but

also leads to a wiring and cross-talk challenge in selectively addressing each qubit. In 2022, single- and two-qubit gates reaching 99% fidelities were demonstrated (21–23) and recently replicated with two-qubit devices fabricated in a 300-mm foundry environment (24). Electron spin qubits on dopants can also be used to facilitate coupling between nuclear spins, with >99% gate fidelities between three nuclear spin qubits reported (25). In Ge/SiGe quantum dots, a relatively new platform (26), universal control of four qubits was recently demonstrated (27).

A key materials challenge for Si/SiGe and Ge/SiGe qubits is achieving uniform and large valley splitting, given that the qubit states can mix with the excited valley states, causing leakage and control imperfections. Another challenge is temperature, as quantum dot spin qubits often operate at temperatures on the order of 100 mK. Although not a showstopper—superconducting qubits are operated at even lower temperatures—there is an increasingly successful push to increase this temperature in both silicon and germanium, with Clifford gate fidelities of 99.8% achieved at a temperature of 1 K (28).

Separate from the above, quantum dots based on III-V heterostructures are being developed toward high-efficiency photon sources (29, 30). Such sources could serve as enabling technology in all four domains and set the quantum dot TRLs in Fig. 1 for networking and sensing.

Spin defects

Individual trapped atoms or ions in solids provide a combination of efficient photonic interfaces [up to about 50% photon detection efficiency at the system level (31)], coherence times up to seconds for electronic spins and minutes for nuclear spins (32, 33), and solid-state hosts that are amenable to on-chip integration with semiconductor electrical and photonic circuits (34). In a quantum network setting, modules contain optically interfaced “communication” qubits such as color center electrons or ions (see next subsection) with nearby nuclear spins or different ions forming a local data qubit register. These systems combine local information processing with long-range optical photon-mediated entanglement links to achieve quantum connectivity for modular quantum computing and quantum internet applications. Additionally, these optically interfaced spins in semiconductors are also highly promising systems for sensing with nanoscale resolution (35).

Networks operating up to five qubits over up to three nodes have been built in the lab using diamond (36, 37) as well as ytterbium-171 ions in yttrium orthovanadate crystals (38); silicon and silicon-carbide also hold promise (39, 40). Employing quantum-frequency conversion to low-loss telecom bands and fiber stabilization architectures have enabled heralded entanglement over tens of kilometers of deployed telecom fiber (41) and between 10-km distant nodes (42). In-lab experiments have shown proof-of-principle demonstrations of quantum communication advantages (31), quantum network primitives (36, 37, 43), and

software-controlled entanglement distribution (44). In these works, single-qubit (two-qubit) logic operates with around 99% (98%) fidelity, whereas the 99.999% (99.93%) logic fidelities recently shown in an optimized two-qubit register (45) indicate the potential for beyond-threshold operation. Gate times on electronic spin qubits are submicro-second, with nuclear spin qubit operation typically a factor of 1000 slower. Within a single-qubit register, mapping of a 50-spin qubit network for simulations (46) has been reported. Entanglement over optical channels currently reaches fidelities from 80 to 90% with rates up to tens of hertz (37, 43).

The main challenges for scaling from these initial demonstrations are systematically increasing operating fidelities and efficiency as well as significantly increasing entanglement distribution rates. For the latter, massive multiplexing enabled by on-chip integration of large numbers of modules is likely required, leading to challenges that are common among all platforms (yield, power, calibration, and cost per qubit).

Atom arrays and trapped ions

Both trapped ions as well as neutral atoms have a long history in quantum technology. One of the earliest two-qubit gates was performed using trapped ions in 1995 (47), and quantum phase transitions were observed in 2002 in ultracold atoms (48). Furthermore, both systems find applications as quantum sensors and are being used as the most precise clocks (49). The two architectures share many commonalities: Qubit states are encoded in energy levels of the atom (typically spin levels, orbital levels, or a combination of both) and are manipulated using laser and microwave fields. The experiments take place in ultrahigh vacuum chambers, which essentially remove the ambient environment and thereby many decoherence sources. Consequently, atomic qubits offer coherence times that extend well beyond 10 s. Because atoms (ions or neutral) are all indistinguishable among the same species, these architectures are naturally homogeneous, which can be advantageous for scaling to multiple identical modules. The largest trapped-ion systems have hundreds of ions (50, 51) and thousands of neutral atoms (52). Both architectures have recently implemented quantum algorithms with logically encoded qubits and gates (53, 54).

A main difference between the two architectures is the way that the individual qubits are being trapped and the two-qubit gates are being implemented. The charge of the ion provides a convenient means to trap it using radio frequency electromagnetic fields in a quadrupole configuration. In this approach, the ions experience the same trapping potential while also experiencing a Coulomb interaction between each other. Two-qubit gates are performed by coupling the common motional modes of the ions to their internal states. Although this results in high-fidelity gates, the scaling beyond 100 qubits within the same trapping potential is challenging. One possible way forward is the quantum charge-coupled device architecture in which ions are moved along a surface trap chip between multiple separate trapping zones (55, 56).

For neutral atoms, trapping potentials are provided by optical fields either in the form of counterpropagating beams that form an optical lattice or by tightly focused laser beams (optical tweezers). The first approach has been very effective in the context of quantum simulation, where atoms tunneling between different potential wells in the lattice are akin to electrons hopping in a solid-state lattice (57). Optical tweezers have recently also advanced quantum information processing for neutral atoms (58). Here, each tweezer hosts an individual qubit, and two-qubit gates are implemented by exciting to high-principal quantum number states (Rydberg states) that lead to strong interactions over typical tweezer distances of a few micrometers. An advantage of this approach is that scaling becomes a question of how many optical tweezers can be produced in the field of view of a microscope, with limits estimated to be around 100,000. Scaling beyond these limits will likely be enabled by a modular architecture, in which optical interfaces will provide the links between separate computation modules (59). In

this direction, prototype quantum networks have been implemented both with trapped ions (60) and neutral atoms (61).

Optical photonic qubits

Current all-photonic quantum computing efforts are primarily focused on cluster-state (one-way) quantum computing (62, 63) and are limited mostly by losses resulting from imperfect photonic devices (see “The challenges” section) as well as inefficient and imperfect quantum light sources based on bulk optical nonlinearities. To achieve improvements, alternative types of sources are being pursued, including high-efficiency quantum dot single-photon sources (29, 30) and those based on cavity quantum electrodynamics (QED) (64). Cavity QED has been used to generate record-size photonic cluster states—12 photons in one dimension (64)—and the extension of this approach based on two atoms strongly coupled to a cavity has been used to perform fusion of two linear clusters into a 2D cluster (65). At present, this approach scales better in success probability than does the all-photonic approach, as it has a high efficiency of single-photon generation and detection, and uses quantum memory (atoms), which acts as an entangler between photons (64).

Irrespective of the quantum technology platform, photonics will play an essential role in scaling of quantum circuits. Photons are ideal for interconnecting quantum systems and for multiplexing control signals for many qubits (66). Photons are also necessary for trapping atoms in tweezer array quantum computers and for controlling and measuring a variety of optically interfaced qubits—atoms, ions, and artificial atoms (color centers in solids, quantum dots). However, despite major progress over the past couple of decades (see Fig. 3), state-of-the-art photonics is lossy and lacks some of the devices and functionalities needed for quantum systems, as discussed in the following section.

The challenges

The technologies in Fig. 1 have all demonstrated the fundamental DiVincenzo criteria (67)—the minimum requirements for a technology to be extensible in principle—with coherent operability at the level of a few qubits to hundreds of qubits. However, realizing the full promise of quantum technologies will require operability at much larger numbers of qubits. Scaling within current monolithic frameworks can support marginal increases in qubit counts, but to truly scale—as with all complex systems—architectural modularity will become necessary. The challenges to scaling are many, and we choose here to focus on four that, in broad terms, are shared by all of these platforms: materials and fabrication, wiring, power, and calibration and control. All of these challenges influence the choices encountered when defining a quantum information processing module that can be interconnected to form larger quantum systems.

Fabrication and materials

Materials science and fabrication engineering are critical to the success of all qubit platforms, whether that is manufacturing the qubits themselves or the ancillary electronics and optics needed to isolate, control, and read them out. The premium is on robust, reproducible, and scalable manufacturing. From an engineering standpoint, it is desirable to leverage existing semiconductor manufacturing while introducing as few new processes as possible. At the same time, quantum technologies are different and will require the targeted introduction of new materials systems, new fabrication flows, and new tooling to realize the advanced performance required.

All technologies in Fig. 1 leverage industry-standard fabrication tools whenever feasible to manufacture qubits and/or their control electronics. For example, superconducting qubits use aluminum, refractory metals, and refractory nitrides on silicon substrates with 3D heterogeneous integration (68), all of which are familiar to silicon fabrication (69). Surface-trapped ions (70) and photonic platforms (62, 63) use integrated optics (e.g., Si/SiN waveguides, optical switches, phase shifters),

cryogenic CMOS control (71), and superconducting photodetectors (72). Spin defects strive for precision lateral and depth placement (73). The challenge lies in reproducible, robust, cost-effective fabrication with sufficiently low loss and low defect density to enable large-scale operation.

Materials science is an important, but not independent, part of the fabrication challenge. When required, new materials that offer enhanced performance, for example, electro-optic barium titanate (62, 74) or strontium titanate (75, 76), must be integrated into a fabrication process flow. The act of patterning, etching, and stripping resist from an otherwise pristine material certainly alters the material surface and possibly its performance. Surface and interface science is particularly important in the context of decoherence mechanisms in superconducting (77), semiconducting (10), spin defects (78), photonics, and trapped-ion qubits (70).

Photonics exemplifies the challenge: Silicon photonics fabricated in commercial foundries typically exhibits propagation loss of 0.25 to 1 dB/cm, and fiber-to-chip grating couplers available in their process design kits exhibit >3 dB of loss and are narrowband (69). Although lower propagation losses of decibels per meter have been reported in research-grade silicon nitride, lithium niobate, or silicon carbide photonics (73, 79, 80), these processes have not yet become widely available in a foundry environment. Moreover, state-of-the-art spatial light modulators have limited resolution and speed for scaling tweezer arrays (81), and solutions are needed at shorter wavelengths, with smaller pixel sizes and higher speed for trapping more atoms, reducing their separation, and controlling and moving atoms faster.

Similar challenges exist in microwave photonics, where deposited dielectrics such as silicon oxide, silicon nitride, and amorphous silicon are commonly used in foundry fabrication flows but exhibit quality factors of 10^3 to 10^5 in the microwave regime (82). Although this may be sufficient for most classical applications, it is generally inadequate for quantum technologies, which even today require quality factors exceeding 10^6 (9). The best quantum transducers and quantum frequency converters today are still very inefficient (41, 83), which impedes their use of photonics in networking microwave superconducting and semiconducting quantum circuits and limits scaling of quantum networks, requiring telecom fiber communication over long distances.

All of these challenges will require materials and fabrication advances in a foundry setting. Academia is well suited for the identification of new candidate materials and fabrication concepts, as students can try a wide range of approaches, fail fast, and prototype successful options. Industry is well suited to taking these concepts, weighing the pros and cons of the often-conflicting options, and then optimizing for robust and reproducible system performance at scale.

Wiring

Once modules of qubits are assembled, the control and readout signals—the input and output of quantum computers—must be routed at scale. Most qubit modalities today require one to three channels per qubit, for the handling (trapped ions, neutral atoms), control, and readout signals, whether realized in the direct current/baseband, microwave, or optical domains. Brute-force scaling of waveguides or wires to the millions of qubits level will become impractical—both in cost and defect density—and alternative approaches must be sought.

Multiplexing signals is one approach to extensible control, where a single channel feeds multiple qubits. As an early example, two-qubit gates in an array of about 300 neutral atom qubits have been implemented by a single laser and coherently moving the atoms in and out of the beam (53). Challenges here include increasing the number of qubits per channel and accommodating the trade between large numbers of multiplexed qubits and the degree of control parallelization that is achievable (simultaneous versus sequential control operations). Another approach is to integrate control electronics with the qubits, whether through (cryogenic) CMOS (21, 71), superconducting logic

(84), integrated optics (70), or even room temperature variants where possible. The challenge—owing to noise, heating, or the trade-off between flexibility and functionality—is in maintaining high-fidelity quantum operations in the presence of these integrated, colocated control electronics and optics. An additional challenge is the development of quantum interconnects that coherently connect modules with sufficiently high fidelity to distribute entanglement throughout the quantum computer (66).

Calibration and control

A third challenge is the calibration of the qubits within and across modules. All qubits require some degree of calibration. Artificial atoms are manufactured qubits, each of which is nominally identical by design but practically different owing to small fabrication variations. Neutral atoms and ions have identical internal degrees of freedom, generally translating to identical single qubits, but the two-qubit operations are susceptible to variations in the trap potentials or the optical lattices that hold them (70). All qubits are susceptible to variations in the control and readout signals. Finally, calibrating qubits and their cross-talk with other qubits becomes a daunting task as the numbers of qubits increase (12). Machine learning is already being applied to assist with this challenge.

Size and power

Size and power are also challenges to extensibility. Today's development focuses largely on achieving operability, without many constraints on size or power. However, beyond-brute-force, large-scale machines will be constrained by both their size and power requirements. For example, superconducting qubits today are quite large, at $\sim 1 \text{ mm}^2$ per qubit (9), which improves performance at the expense of real estate. Many of the technology platforms require large-scale cryogenic systems to cool the computer (qubits) and/or its peripherals (e.g., photodetectors, control electronics) (9, 70). Neutral atom quantum computers rely on high-power lasers (53). As these technologies evolve, both the size and power will increasingly constrain the trade space and necessitate innovation to keep the cost and footprint manageable.

Outlook: What is going to give us the 1000× boost?

Modularity and architecture

Zooming out on the challenges outlined in the previous section, it is clear that although improvements can be made by optimizing elements within monolithic qubit systems, looking for 1000× improvements forces us to look at different architectures. In particular, building modularity into the systems provides a pathway that breaks the scaling into more manageable pieces. The key idea is to assemble full systems from a large set of smaller modules (Fig. 4) that are interconnected by quantum links (85). Such modularity will cap the wiring, size, (cooling) power, control, and calibration challenges to those of a single module. Note that modularity may be realized at different levels (e.g., within a chip by using on-chip wiring to connect small architectural modules, or between chips and modules using 3D heterogeneous integration or optical interconnects).

Besides removing scaling barriers, modularity opens the possibility of increased as well as dynamic connectivity throughout the system, enabling new functionality. One example is the use of different types of error-correction codes on quantum computers such as qLDPC (quantum low-density parity check) codes that could reduce the overall demands on the system (e.g., lower qubit number overhead). In addition, modular approaches may optimize different functionality in different parts of the system. For instance, quantum memory units might be separately optimized from quantum processing units.

Although modular architectures can thus help solve the scaling challenges outlined above and enable new opportunities, this approach comes with a major challenge: How can one efficiently interconnect these modules while maintaining the required performance?




Fig. 4. Future modular architecture. Each quantum computing module (inset) consists of submodules with different functionalities as well as a classical control layer. These modules are then connected to each other through short quantum interconnects resembling a rack architecture in a data center. These racks can then be connected to others across longer distances in a quantum internet. QRAM, quantum random access memory.

The precise quantitative requirements for quantum interconnects (66) will differ by platform and specific application, but, qualitatively, they encompass fidelity, speed, and availability of the interconnect link. Heralding and error detection procedures allow for high-fidelity entanglement to be shared between modules but are intrinsically probabilistic. Direct transmission of qubits between chips (e.g., through transduction) is in principle deterministic, but unavoidable loss will also make these procedures probabilistic in practice. This probabilistic nature can be remedied, for instance, by adding long-lived quantum memories to make the links near-deterministic or by accounting for them in the error correction codes directly.

Heterogeneous systems

The need for modularity and integration will require heterogeneous systems at multiple levels of the architecture. For instance, the different modules for interconnecting separate computing nodes (Fig. 4) will likely consist of different modalities than the processor itself. Therefore, hybrid interfaces will need to be developed that enable quantum states to be shared between two or more different subsystems (66). Depending on the qubit implementation, these interfaces will need to bridge large frequency regimes, ranging from microwave to optical. The development of such interfaces will allow for the integration of different subsystems that are optimized for different quantum information tasks such as quantum memory, processing, and communication. Beyond heterogeneous quantum systems, tight integration of the classical and quantum parts of information processors will become increasingly important (Fig. 4).

Integration and interfaces

Photonics is key for scaling almost any type of quantum system, because of its essential role in quantum interconnects and multiplexing of control signals. But how can one achieve a giant boost in photonics performance needed for its inclusion into functional quantum systems? Toward this goal, innovative photonics design approaches are emerging, driven both by quantum technologies and the need for better optical interconnects for artificial intelligence (AI) hardware (86). New and improved photonic materials are also emerging to implement functionalities necessary for scaling quantum circuits. This includes thin-film lithium niobate for switching and microwave-to-optical transduction devices (79); more-powerful cryogenic electro-optic materials such as strontium titanate (75, 76); wider bandgap and lower-loss materials such as silicon carbide, silicon nitride, and aluminum nitride (72, 80, 87); materials that host rare earth ions for quantum networking (38, 88); and even laser systems on chip (89). Flat optics and metasurfaces (90) are candidates for scaling of tweezer arrays and atom control circuits, in combination with integrated photonics (91). In the microwave domain, low-loss waveguides and couplers are being 3D-integrated to distribute entanglement among chiplets within and between modules (9).

In all-photonic quantum computing, the efforts to employ high-efficiency single-photon sources are ongoing (29, 30). In addition, cavity QED approaches to cluster state generation, which have been used to generate record-size clusters (64, 65) could be translated from AMO to solid-state systems (73), enabling further scaling relative to atomic cavity QED platforms.

Conclusions

The past decade has witnessed extraordinary progress in the development of new hardware for quantum information processing across many platforms. This rapid acceleration has, from the start, been driven by collaborations spanning academia, industry, and national facilities. These assemblies of talent and infrastructure enable researchers to focus on material and device challenges unlike any we have experienced, from deterministic atomic-scale control of matter to the engineering of quantum entanglement. Continuing progress also hinges on advancing current technologies that will be required to deploy hybrid devices. Although quantum technologies have moved to higher TRLs, breakthroughs in performance are needed across all platforms to unlock the full power of quantum computing, entanglement-enhanced sensing, and a global quantum internet. These are extraordinary challenges that rise beyond any individual group, company, or nation. Such challenges continue to be surmounted through a combination of foundational scientific research and innovative engineering. Moreover, this is a global enterprise that is raising important ethical and legal questions ranging from privacy to the establishment of new quantum standards. Finally, the future of this field may largely depend on establishing a new legion of quantum scientists and engineers who will deliver a new form of information processing and transformative applications to society.

REFERENCES AND NOTES

1. AI Impacts, AI Timelines: Trends in the cost of computing, 10 March 2015; <https://aiimpacts.org/trends-in-the-cost-of-computing/>.
2. C. W. Gwyn, R. Stulen, D. Sweeney, D. Attwood, *J. Vac. Sci. Technol. B* **16**, 3142–3149 (1998).
3. Q. Xu, B. Schmidt, S. Pradhan, M. Lipson, *Nature* **435**, 325–327 (2005).
4. C. Gidney, M. Ekerå, *Quantum (Vienna)* **5**, 433 (2021).
5. C. Gidney, arXiv:2505.15917 [quant-ph] (2025).
6. A. Caesura *et al.*, *PRX Quantum* **6**, 030337 (2025).
7. S. Wehner, D. Elkouss, R. Hanson, *Science* **362**, eaam9288 (2018).
8. K. Azuma *et al.*, *Rev. Mod. Phys.* **95**, 045006 (2023).
9. M. Kjaergaard *et al.*, *Annu. Rev. Condens. Matter Phys.* **11**, 369–395 (2020).
10. G. Burkard, T. D. Ladd, A. Pan, J. M. Nichol, J. R. Petta, *Rev. Mod. Phys.* **95**, 025003 (2023).
11. Y. Kim *et al.*, *Nature* **618**, 500–505 (2023).
12. Google Quantum AI and Collaborators, *Nature* **638**, 920–926 (2025).
13. D. Rower *et al.*, *PRX Quantum* **5**, 040342 (2024).

14. F. Marxer *et al.*, arXiv:2508.16437 [quant-ph] (2025).
15. Google Quantum AI, *Nature* **614**, 676–681 (2023).
16. Y. Sunada *et al.*, *Phys. Rev. Appl.* **17**, 044016 (2022).
17. F. Swiadek *et al.*, *PRX Quantum* **5**, 040326 (2024).
18. S. Krinner *et al.*, *Nature* **605**, 669–674 (2022).
19. S. Krinner *et al.*, *EPJ Quantum Technol.* **6**, 2 (2019).
20. S. Storz *et al.*, *Phys. Rev. Lett.* **135**, 030801 (2025).
21. X. Xue *et al.*, *Nature* **601**, 343–347 (2022).
22. A. R. Mills *et al.*, *Sci. Adv.* **8**, eabn5130 (2022).
23. A. Noiri *et al.*, *Nature* **601**, 338–342 (2022).
24. P. Steinacker *et al.*, *Nature* **646**, 81–87 (2025).
25. I. Thorvaldson *et al.*, *Nat. Nanotechnol.* **20**, 472–477 (2025).
26. G. Scappucci *et al.*, *Nat. Rev. Mater.* **6**, 926–943 (2021).
27. X. Zhang *et al.*, *Nat. Nanotechnol.* **20**, 209–215 (2025).
28. J. Y. Huang *et al.*, *Nature* **627**, 772–777 (2024).
29. P. Senellart, G. Solomon, A. White, *Nat. Nanotechnol.* **12**, 1026–1039 (2017).
30. M. L. Chan, A. A. Capatos, P. Lodahl, A. Søndberg Sørensen, S. Paesani, arXiv:2507.16152 [quant-ph] (2025).
31. M. K. Bhaskar *et al.*, *Nature* **580**, 60–64 (2020).
32. C. E. Bradley *et al.*, *Phys. Rev. X* **9**, 031001 (2019).
33. C. P. Anderson *et al.*, *Sci. Adv.* **8**, eabm5912 (2022).
34. L. Li *et al.*, *Nature* **630**, 70–76 (2024).
35. C. L. Degen, F. Reinhard, P. Cappellaro, *Rev. Mod. Phys.* **89**, 035002 (2017).
36. M. Pompili *et al.*, *Science* **372**, 259–264 (2021).
37. Y.-C. Wei *et al.*, *Science* **388**, 509–513 (2025).
38. A. Ruskuc *et al.*, *Nature* **639**, 54–59 (2025).
39. S. Simmons, *PRX Quantum* **5**, 010102 (2024).
40. N. T. Son *et al.*, *Appl. Phys. Lett.* **116**, 190501 (2020).
41. C. M. Knaut *et al.*, *Nature* **629**, 573–578 (2024).
42. A. J. Stolk *et al.*, *Sci. Adv.* **10**, eadp6442 (2024).
43. S. L. N. Hermans *et al.*, *Nature* **605**, 663–668 (2022).
44. C. Delle Donne *et al.*, *Nature* **639**, 321–328 (2025).
45. H. P. Bartling *et al.*, *Phys. Rev. Appl.* **23**, 034052 (2025).
46. G. L. van de Stolpe *et al.*, *Nat. Commun.* **15**, 2006 (2024).
47. C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, D. J. Wineland, *Phys. Rev. Lett.* **75**, 4714–4717 (1995).
48. M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, I. Bloch, *Nature* **415**, 39–44 (2002).
49. A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, P. O. Schmidt, *Rev. Mod. Phys.* **87**, 637–701 (2015).
50. D. Kiesenhofer *et al.*, *PRX Quantum* **4**, 020317 (2023).
51. S.-A. Guo *et al.*, *Nature* **630**, 613–618 (2024).
52. H. J. Manetsch *et al.*, *Nature* **647**, 60–67 (2025).
53. D. Bluvstein *et al.*, *Nature* **626**, 58–65 (2024).
54. C. Ryan-Anderson *et al.*, arXiv:2404.16728 [quant-ph] (2024).
55. D. Kielpinski, C. Monroe, D. J. Wineland, *Nature* **417**, 709–711 (2002).
56. J. M. Pino *et al.*, *Nature* **592**, 209–213 (2021).
57. C. Gross, I. Bloch, *Science* **357**, 995–1001 (2017).
58. M. Morgado, S. Whitlock, *AVS Quantum Sci.* **3**, 023501 (2021).
59. J. Covey, H. Weinfurter, H. Bernien, *NPJ Quantum Inf.* **9**, 90 (2023).
60. V. Krutyanskiy *et al.*, *Phys. Rev. Lett.* **130**, 050803 (2023).
61. W. Zhang *et al.*, *Nature* **607**, 687–691 (2022).
62. PsiQuantum team, *Nature* **641**, 876–883 (2025).
63. H. Aghaee Rad *et al.*, *Nature* **638**, 912–919 (2025).
64. P. Thomas, L. Ruscio, O. Morin, G. Rempe, *Nature* **608**, 677–681 (2022).
65. P. Thomas, L. Ruscio, O. Morin, G. Rempe, *Nature* **629**, 567–572 (2024).
66. D. Awschalom *et al.*, *PRX Quantum* **2**, 017002 (2021).
67. D. P. DiVincenzo, arXiv:cond-mat/9612126 [cond-mat.mes-hall] (1996).
68. D. Rosenberg *et al.*, *NPJ Quantum Inf.* **3**, 24 (2017).
69. M. Rakowski *et al.*, 2020 Optical Fiber Communications Conference and Exhibition (OFC), paper T3H.3 (2020).
70. C. Bruzewicz, J. Chiaverini, R. McConnell, J. M. Sage, *Appl. Phys. Rev.* **6**, 021314 (2019).
71. J. Anders *et al.*, *IEEE Trans. Quantum Eng.* **4**, 5100230 (2023).
72. I. Holzman, Y. Imry, *Adv. Quantum Technol.* **2**, 1800058 (2019).
73. D. M. Lukin, M. A. Guidry, J. Vučković, *PRX Quantum* **1**, 020102 (2020).
74. S. Abel *et al.*, *Nat. Mater.* **18**, 42–47 (2019).
75. C. P. Anderson *et al.*, *Science* **390**, 394–399 (2025).
76. A. Ulrich *et al.*, *Science* **390**, 390–393 (2025).
77. I. Siddiqi, *Nat. Rev. Mater.* **6**, 875–891 (2021).
78. L. V. H. Rodgers *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **121**, e2316032121 (2024).
79. D. Assumpcao *et al.*, *Nat. Commun.* **15**, 10459 (2024).
80. C. Xiang, W. Jin, J. E. Bowers, *Photon. Res.* **10**, A82–A96 (2022).
81. R. Lin *et al.*, *Phys. Rev. Lett.* **135**, 060602 (2025).
82. W. D. Oliver, P. B. Welander, *MRS Bull.* **38**, 816–825 (2013).
83. J. Rochman, T. Xie, J. G. Bartholomew, K. C. Schwab, A. Faraon, *Nat. Commun.* **14**, 1153 (2023).
84. N. Takeuchi, T. Yamae, C. L. Ayala, H. Suzuki, N. Yoshikawa, *IEICE Trans. Electron.* **105**, 251–263 (2022).
85. C. R. Monroe, R. J. Schoelkopf, M. D. Lukin, *Sci. Am.* **314**, 50–57 (2016).
86. K. Y. Yang *et al.*, *Nat. Commun.* **13**, 7862 (2022).
87. X. Liu, A. W. Bruch, H. X. Tang, *Adv. Opt. Photonics* **15**, 236–317 (2023).
88. S. Ourari *et al.*, *Nature* **620**, 977–981 (2023).
89. J. Yang *et al.*, *Nature* **630**, 853–859 (2024).
90. A. H. Dorrah, J.-S. Park, A. Palmieri, F. Capasso, *Nat. Commun.* **16**, 3126 (2025).
91. I. Christen *et al.*, *Nat. Commun.* **16**, 82 (2025).

ACKNOWLEDGMENTS

We thank M. Eriksson, S. Guha, and B. Pingault for their assistance with this manuscript. **Funding:** D.D.A. was supported by Q-NEXT, a US Department of Energy Office of Science National Quantum Information Science Research Center, and the Air Force Office of Scientific Research. R.H. was supported by the Dutch Research Council (NWO) through the Spinoza Prize 2019 (project SPI 63-264). J.V. was supported by the Vannevar Bush Faculty Fellowship from the US Department of Defense. H.B. gratefully acknowledges funding from the US Office of Naval Research (grant N00014-23-1-2540) and the US Army Research Office (grant W911NF2410388). W.D.O. gratefully acknowledges funding from the US Army Research Office under award W911NF-23-1-0045. **Competing interests:** D.D.A. is a cofounder and shareholder of Blanqet LLC. W.D.O. was a cofounder and shareholder of Atlantic Quantum, a start-up company focused on superconducting quantum computing. Atlantic Quantum was acquired by Google in October 2025. R.H. is cofounder and shareholder of Delft Networks, a start-up company focused on quantum network products and applications. J.V. is a cofounder and a lead scientific adviser of Spins Photonics Inc, a photonic inverse design software company. H.B. is a cofounder and shareholder of Cavilinq, Inc. **License information:** Copyright © 2025 the authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original US government works. <https://www.science.org/about/science-licenses-journal-article-reuse>

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.adz8659
Supplementary Text

Submitted 8 September 2025; accepted 5 November 2025

10.1126/science.adz8659