IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 12 February 2026, accepted 28 February 2026, date of publication 5 March 2026, date of current version 10 March 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3670749

TOPICAL REVIEW

Comparative Analysis of Contemporary Quantum Computer Processors: Architectures, Performance, and Perspectives

MARKO VUKŠIĆ, JASMIN ĆELIĆ, (Senior Member, IEEE), AND ALEKSANDAR CUCULIĆ
Department of Electrical Engineering, Automation and Computing, Faculty of Maritime Studies, University of Rijeka, 51000 Rijeka, Croatia

Corresponding author: Marko Vukšić (marko.vuksic@uniri.hr)

This work was supported by the Project Line Znanstveno-istraživačka projektna inicijativa (ZIP) University of Rijeka (UNIRI) of the UNIRI for the Research Project ‘‘Sustainable Transformation of Recreational Vessels in the Republic of Croatia’’ under Grant UNIRI-ZIP-2103-8-22.

ABSTRACT Quantum computing promises to revolutionize several scientific and industrial domains by tackling problems that are intractable for classical supercomputers. At the centre of this effort is the quantum processor unit (QPU), implemented today across multiple physical qubit platforms. This review provides a comparative analysis of the state of quantum processors as of late 2024/early 2025, focusing on superconducting qubits, trapped ions, photonic systems, neutral atoms, silicon spin qubits, and emerging topological approaches. We compile and compare recent processor generations such as IBM's Heron family, Google's updated Sycamore-class devices, Quantinuum's H-series, IonQ's Aria and Forte, neutral-atom platforms with more than 1,000 trapped atoms, and silicon spin prototypes like Tunnel Falls. Key performance metrics, including qubit count, gate and readout fidelity, coherence time, operation speed, and holistic benchmarks such as Quantum Volume (QV) and Algorithmic Qubits (#AQ) are examined with reference to concrete experimental milestones. We discuss algorithmic capabilities demonstrated on current hardware, early realizations of logical qubits and fault-tolerant primitives, and the distinction between feasibility-oriented ‘‘quantum advantage’’ experiments and emerging signs of practical ‘‘quantum utility’’. Finally, we outline the main technical bottlenecks for each platform on the path from today's NISQ devices toward universal fault-tolerant quantum computers.

INDEX TERMS Fault-tolerant quantum computing, NISQ era, photonic qubits, quantum computing, quantum processors, quantum volume, qubits, superconducting qubits, trapped ions.

I. INTRODUCTION

Quantum computing is a computing paradigm that uses principles of quantum mechanics; most notably superposition and entanglement to perform calculations beyond the reach of classical computers [1], [2]. Ever since Feynman's early proposals, researchers have anticipated applications in drug discovery, materials science, optimization, finance, cryptography, and fundamental science [2], [3]. Quantum algorithms such as Shor's factoring algorithm and Grover's search algorithm theoretically offer exponential and quadratic speedups, respectively, over the best-known classical algorithms [4]. In practice, however, realizing these advantages requires quantum hardware with many high-quality qubits. In recent years, variational quantum algorithms like the Variational Quantum Eigensolver (VQE) for chemistry and the Quantum Approximate Optimization Algorithm (QAOA) for combinatorial problems have emerged as promising approaches for noisy intermediate-scale devices [5].

Significant progress over the last decade has moved quantum hardware from single-qubit experiments to the current ‘‘Noisy Intermediate-Scale Quantum’’ (NISQ) era [6]. NISQ devices contain tens to a few hundred physical qubits but are still limited by noise and lack full error correction [6]. In 2019, Google achieved a key milestone in quantum computing by executing a random circuit sampling task on its

The associate editor coordinating the review of this manuscript and approving it for publication was Mu-Yen Chen.

53-qubit Sycamore processor; an operation that surpassed the capabilities of classical supercomputers for that specific benchmark [4]. Although the task had limited real-world applicability, it provided the first practical demonstration of quantum advantage on a specific benchmarking task using NISQ hardware. Meanwhile, quantum annealers (e.g., D-Wave systems) and analog quantum simulators have addressed certain optimization and simulation tasks, though this review will focus on gate-model QPUs.

In the last few years, several platforms have demonstrated capabilities that go beyond isolated physics experiments. Superconducting devices have implemented early logical qubits with error suppression by increasing code distance, trapped-ion systems have realized real-time fault-tolerant error correction on small logical memories, and neutral-atom and photonic processors have reported quantum advantage in specific sampling and simulation tasks. Silicon spin qubits, while still operating at smaller scales, have progressed from single- or few-qubit demonstrations to CMOS-compatible prototype arrays. These developments motivate a hardware-level comparison that is explicitly anchored in the pre–fault-tolerant landscape of the mid-2020s rather than a timeless overview of quantum computing.

In this context, it is important to distinguish between quantum advantage and quantum utility. Quantum advantage typically refers to milestone or feasibility demonstrations in which a quantum processor outperforms classical approaches on a well-defined task, often under controlled experimental conditions. In contrast, quantum utility emphasizes the ability to solve practically relevant problems with reproducible and stable performance, even if the advantage over classical methods is modest or domain-specific.

In parallel with hardware development, quantum software has matured considerably. Frameworks such as Qiskit, Cirq, PennyLane, and TKET have empowered researchers to design, simulate, and deploy quantum algorithms with increasing ease. Hybrid quantum-classical workflows are now commonly used in variational algorithms, with classical optimization loops guiding quantum circuits toward optimal configurations. Furthermore, advances in quantum compilers and transpilers have improved fidelity by tailoring circuits to specific hardware topologies and error models. The parallel advancement of both quantum hardware and software plays a critical role in unlocking useful applications during the NISQ era [7].

The global significance of quantum computing has extended far beyond experimental success, prompting robust investments from both governments and private industries. Programs like the U.S. National Quantum Initiative, the European Union's Quantum Flagship, and China's QUESS project have helped fast-track quantum research and infrastructure development [8]. Policymakers now view quantum technologies as essential to national security and economic leadership. These trends have fueled international competition as well as cooperation, reflected in various funding frameworks, collaborative research projects, and institutional alliances. Most national programs now track progress not only through qubit counts, but also through logical-qubit lifetimes, benchmark reproducibility, and application-level demonstrations. These criteria increasingly shape how the performance and maturity of quantum hardware are assessed across different technologies.

At the same time, major tech corporations such as IBM, Google, Microsoft, Intel, Amazon, and Alibaba have developed proprietary quantum platforms, launched public-access quantum services via the cloud, and contributed to open-source toolchains supporting algorithm development and experimentation. For instance, IBM's Quantum System One and roadmap toward fault-tolerant modular architectures, and Microsoft's Azure Quantum platform integrating third-party hardware providers, are examples of how industry accelerates both access and innovation [9]. Cloud-based quantum computing has enabled broader participation from academia and startups, significantly lowering entry barriers for experimentation and development.

The physical realization of qubits plays a foundational role in the development of quantum processing units (QPUs). Various platforms are under active investigation, each offering different balances between coherence time, operational speed, gate fidelity, connectivity, and potential for scaling [10], [11]. Among these, superconducting and trapped-ion qubits currently dominate due to their relative technological maturity and extensive optimization. Photonic qubits stand out for their compatibility with ambient temperatures and built-in communication pathways, while neutral atom systems leverage atomic control to achieve scalable architectures [12], [13]. Silicon-based spin qubits draw on established semiconductor manufacturing methods to enable potential integration. Topological qubits, still largely theoretical, seek to store information in non-local quasiparticles like Majorana modes to intrinsically suppress errors [14].

Table 1 in the following section provides a comparative overview of representative quantum processors across these technologies. The diversity and core characteristics of each platform are also depicted schematically in Figure 1, offering a visual summary of the current qubit technology landscape. Figure 1 focuses on physical qubit technologies and their characteristic trade-offs. Zapata Quantum's Orquestra is included only as a separately labeled software orchestration platform to reflect its role in coordinating hybrid quantum–classical workflows, and not as a qubit technology or hardware QPU implementation.

Each hardware approach reflects a distinct compromise between coherence, fidelity, control complexity, and integration potential; factors that shape both near-term performance and long-term scalability. Recent surveys increasingly distinguish between theoretical capabilities and experimentally verified performance, and the analysis presented here follows this direction by emphasizing system-level benchmarks, platform-specific constraints, and experimentally

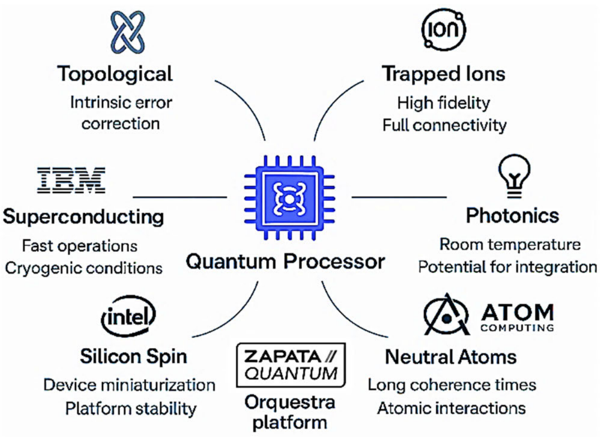

FIGURE 1. **Overview of dominant qubit technologies and their key characteristics, generated using Claude 3.5 (Anthropic, 2025).**

reported results from devices released between 2023 and 2025. Building on this technological landscape, the following section examines the physical implementation of leading qubit modalities alongside key performance metrics and engineering challenges.

II. PHYSICAL REALIZATIONS OF QUBITS AND LEADING PLATFORMS

Several qubit technologies currently dominate quantum computing research and industrial development [10], [11]. These platforms utilize radically different physical systems, ranging from superconducting circuits and trapped ions to photonic systems and neutral atoms, to represent and manipulate qubits. Each implementation comes with distinct engineering challenges, operational environments, and performance trade-offs.

The implementation of qubits across various quantum systems shapes not only device architecture but also the nature of computational workflows and algorithm design. Instead of a single dominant standard, the field has evolved into a landscape of competing approaches, each defined by different trade-offs between coherence times, connectivity schemes, gate speeds, fabrication maturity, and integration prospects.

Rather than converge on one solution, research and industry have diversified: superconducting circuits advance rapidly due to compatibility with microfabrication techniques, while atomic-scale platforms such as trapped ions or neutral atoms prioritize coherence and flexible interactions. Some directions, like topological qubits, are still speculative but hold promise for long-term error resilience. Because each qubit platform presents its own unique set of strengths and limitations, developers have begun placing greater emphasis on more comprehensive performance assessments [15]. Instead of relying solely on qubit numbers, newer benchmarking methods evaluate system-level capabilities under practical algorithmic workloads. These refined metrics aim to gauge a quantum system's actual utility in solving meaningful computational problems. Concepts like Quantum Volume or Algorithmic Qubits are now used to reflect practical computational power more accurately [15].

Major research institutions and private companies alike continue to push development forward through distinct technological strategies. Industry leaders such as IBM, Google, IonQ, and Quantinuum are actively advancing qubit modalities and architectures that align with their respective visions for building scalable and reliable quantum systems [9].

Table 1 below presents a comparative summary of representative quantum processors, covering qubit type, scale, performance benchmarks, and key engineering challenges. This serves as a basis for the subsequent detailed discussion of each qubit technology and the platforms that implement them.

From a practical perspective, current quantum platforms can be broadly grouped according to their near-term NISQ applicability and their long-term potential for fault-tolerant quantum computing. Superconducting and trapped-ion systems currently lead in short-term NISQ applications due to their relatively mature control, cloud accessibility, and demonstrated performance in hybrid quantum–classical workflows. Neutral-atom and photonic architectures, while still developing robust gate fidelities and error mitigation strategies, offer strong long-term scalability prospects through flexible connectivity and reconfigurable architectures. Silicon spin qubits represent a promising bridge between near-term and long-term approaches through compatibility with semiconductor manufacturing, whereas topological qubits remain primarily a long-term research direction targeting intrinsic error resilience. This layered landscape suggests that near-term quantum utility and long-term fault tolerance may emerge from different technological pathways.

Before examining each hardware platform, it's helpful to understand that no single qubit technology holds all the answers. Some are already running basic quantum algorithms and showing promise in specialized areas, while others remain experimental, with researchers still exploring their practical potential. Each approach reflects a different trade-off, from how qubits are built and controlled to how well they might scale in the future. In the following sections, we present a deeper examination of the leading qubit modalities; superconducting qubits, trapped ions, photonic qubits, neutral atoms, silicon spin qubits, and topological qubits, highlighting the physics behind their operation, recent performance milestones, and the challenges they must overcome to scale toward fault-tolerant quantum computation.

Table 1 should be interpreted as a high-level comparative snapshot rather than a definitive ranking, since different architectures are optimized for distinct objectives. The detailed discussion in the following subsections builds directly on this overview, linking hardware characteristics to algorithmic capabilities and engineering trade-offs. This distinction is particularly important because different platforms are optimized for either short-term NISQ performance or long-term fault-tolerant scalability.

TABLE 1. ***(Continued.)*** **Comparative overview of contemporary quantum processors and their characteristics.**

Company /Institution	PROCESSOR (LATEST/SIGNIFICANT)	QUBIT TYPE	NUMBER OF QUBITS (REPORTED)	KEY FEATURES /FOCUS	BENCHMARKS (E.G., QV, #AQ)	Main Technology Challenges
IBM Quantum	Heron (133 qubits)	Superconducting	133	Low error rate (3x better than Eagle), modularity, new connectors (Kookaburra) for chip linking. Focus on quality over mere number.	High Quantum Volume (QV) targeted.	Cooling, noise reduction, scaling with quality.
Google Quantum AI	Sycamore (upgrades)	Superconducting	53-70+	Quantum supremacy demonstration, algorithm research, fidelity improvement, and error reduction.	Focus on specific problems (e.g., material simulation).	Error rate reduction, scaling, long-term coherence.
IonQ	Forte	Trapped Ions	32	High coherence and fidelity, full all-to-all qubit connectivity, dynamic reconfiguration.	#AQ 29 (Algorithmic Qubits)	Gate speed, scaling ion numbers while maintaining control.
Quantinuum	H2	Trapped Ions	32	High fidelity, low error rates, "race-track" architecture, TKET integration.	High QV (e.g., 8192 for H1-1, H2 targets higher), #AQ.	Gate speed, scaling "race-track" architecture.
Rigetti Computing	Aspen-M-3 (84 qubits)	Superconducting	84	Modular architecture, fast gates.	Focus on improving QV.	Inter-module connectivity, error reduction.
Intel	Tunnel Falls (12 qubits)	Silicon Spin	12	Fabrication on 300mm CMOS wafers, potential for mass production. High qubit uniformity.	Developmental stage, focus on foundational technology.	Increasing qubit count, improving coherence and fidelity.
Atom Computing	Phoenix (prototype)	Neutral Atoms	1000+ (atoms)	Large number of atoms, potential for scaling, dynamic atom rearrangement.	Early stage, focus on number scaling and control.	Qubit quality from a large atom number, two-qubit gate speed.

TABLE 1. Comparative overview of contemporary quantum processors and their characteristics.

PsiQuantum	(No named "processor")	Photonic	Target: 1 million+	Long-term goal: fault-tolerant computer.	N/A (no public benchmarks of small systems yet)	Photon generation & detection, losses, deterministic gates.
Xanadu	Borealis (216 qubits)	Photonic (squeezed light)	216 (squeezed modes)	Quantum advantage on GBS problem. Programmable.	Specific tasks (GBS).	Universality, implementing more complex quantum gates.
Microsoft Quantum	(No "processor")	Topological (Majorana)	Research (target: 1+)	Fundamental research, goal: inherently fault-tolerant qubits.	N/A (fundamental research phase)	Creating & controlling Majorana fermions, scientific proof.

A. SUPERCONDUCTING QUBITS

Superconducting qubits, most commonly implemented as transmons, circuit elements incorporating Josephson junctions, represent one of the most technologically developed approaches in quantum computing today [6]. These circuits operate at extremely low temperatures (millikelvin range), where electrical resistance and thermal noise are effectively suppressed. Over the past twenty years, their performance has markedly advanced, with no inherent physical limitations identified thus far [16].

In recent devices, the spread in qubit frequencies, junction variability, and residual two-level system defects remain the dominant sources of coherence limitations. While these effects are well documented, they continue to influence gate scheduling, spectator-qubit errors, and crosstalk on chips above roughly 100 qubits.

Industry leaders such as IBM and Google have been instrumental in pushing this technology forward. IBM's devices like the 127-qubit Eagle and the 433-qubit Osprey utilize fixed-frequency transmons arranged in a heavy-hexagon topology. This layout provides a practical compromise between minimizing crosstalk and maintaining useful connectivity; each qubit typically interacts with a limited number of neighbors [7], [17], [18]. The newer 133-qubit Heron processor emphasizes high-fidelity gate operations (>99.9% single-qubit, ∼99.5% two-qubit) under optimized calibration and benchmarking conditions and is designed with modular scalability in mind [14], [15]. Google's Sycamore chip, initially with 53 qubits, achieved ∼99.4% two-qubit fidelity in experimental demonstrations and was central in the first demonstration of quantum advantage through random circuit sampling [5].

In parallel with these U.S.-based platforms, European efforts have also advanced, most notably those led by IQM, which develops superconducting processors based on a tunable-coupler architecture [8]. IQM's devices, deployed in several national quantum initiatives across Europe, emphasize stability of two-qubit gates and system modularity rather than rapid increases in qubit count. Recent installations, typically in the 20–50 qubit range, have been used to explore application-oriented workloads in chemistry and optimization. Although these systems do not yet operate at the scale of IBM's or Google's flagship processors, their focus on integrated control electronics and cryogenic co-design addresses practical engineering considerations that increasingly shape real-world performance of superconducting platforms.

Modern superconducting platforms support nanosecond-scale operations (typically 10–100 ns), enabling fast gate execution [19], [20]. These systems also benefit from compatibility with standard microwave control and readout hardware, leveraging technologies developed in classical integrated circuits. Entanglement is typically achieved using microwave-driven ZZ or ZX interactions. These hardware capabilities are further supported by mature control-software stacks, such as IBM's Qiskit, which provide automated calibration tools and hardware-aware circuit transpilation.

Superconducting qubit processors have been employed to run simplified instances of known algorithms. For example, IBM demonstrated Shor's algorithm factoring 21 on a 5-qubit transmon device, and Grover's search algorithm has been tested using error mitigation on a small subset of a 127-qubit system [21], [22], [23]. These experimental results validate hardware operation, though full-scale versions of such algorithms will necessitate significant error correction and larger qubit counts.

Nonetheless, several engineering obstacles remain. Maintaining cryogenic environments imposes technical and financial complexity, especially when considering systems with thousands of qubits. Enhancing coherence time, currently around 100 μs on average, demands further improvements in materials, particularly better dielectric interfaces and 3D integration strategies [9], [24]. Another concern is limited on-chip connectivity: transmons in planar arrays typically connect to only 2–4 neighbors, which increases circuit depth due to required SWAP operations for distant interactions.

A growing distinction exists between laboratory-grade demonstrations and cloud-deployed systems operating continuously. While state-of-the-art devices can exceed 99% two-qubit fidelity under optimized calibration, long-term stability and cross-device reproducibility remain active engineering challenges.

Despite these constraints, progress has been consistent. IBM, for instance, reported Quantum Volume (QV) milestones of $2^8=256$ in 2019 and $2^{10}=1024$ in 2020, with a roadmap goal of reaching 2^{20} by 2025 [15], [25], [26]. Given their blend of rapid gate execution, manufacturing compatibility, and extensive software infrastructure, superconducting qubits remain a leading candidate for scalable quantum computing; provided ongoing engineering challenges can be successfully addressed.

B. TRAPPED IONS

Quantum processors based on trapped ions utilize individual charged atoms, held in place using electromagnetic fields, to serve as qubits. These qubit states are typically defined between two hyperfine levels of the ion's ground electronic state, offering remarkable stability. In some cases, coherence times exceeding several minutes have been recorded [27]. Operations are carried out by directing laser beams or microwave fields at the ions to manipulate their collective vibrational modes, which enables entangling gates between qubits [28].

One of the standout benefits of this approach lies in its gate fidelity: single-qubit operations often reach fidelities above 99.9%, while two-qubit gates can achieve between 99.5% and 99.9% in controlled laboratory environments [10], [29]. Additionally, trapped ions are indistinguishable and can be collectively cooled and addressed, enabling near all-to-all connectivity within the same trap. Entangling gates exploit the shared motional modes of the ion chain, which allows long-range interactions without requiring physical proximity. These physical advantages have translated reasonably well to commercial systems, although practical implementations introduce additional constraints compared with controlled laboratory setups.

Commercial progress in this domain is exemplified by platforms such as IonQ and Quantinuum (formerly part of Honeywell). IonQ's systems, built using Ytterbium ions arranged in linear chains, have scaled up to 32 qubits, maintaining high gate fidelity throughout. Their Aria and Forte devices report Algorithmic Qubit (AQ) counts of 20 and 29, respectively, signifying the ability to execute deep, error-mitigated circuits involving meaningful quantum workloads [30]. These platforms have been employed in experimental demonstrations of quantum error correction codes and variational chemistry algorithms.

Quantinuum's H-series processors, such as H1 and H2, adopt a racetrack-style trap to shuttle Yb/Sr ions between separate zones. This dynamic architecture reduces unwanted crosstalk by isolating ions during critical operations [29]. In 2021, the H1-1 processor, hosting 20 qubits, achieved a record Quantum Volume of 8192, reflecting both the fidelity and utility of the system [31]. Historically, trapped-ion devices were among the first to demonstrate core algorithms like quantum teleportation and basic instances of Shor's factoring, such as factoring 15, in academic research settings.

However, scalability and operational speed remain primary challenges. Two-qubit gates in trapped-ion systems typically operate on a timescale of tens to hundreds of microseconds; significantly slower than in solid-state qubit platforms [32]. This latency stems from the reliance on collective motion of ions, which becomes increasingly complex as more ions are added. As the number of ions in a trap increases, their collective vibrational modes become more closely spaced, which makes precise gate operation more difficult and can degrade performance. To mitigate these effects, scientists are developing modular approaches, such as linking smaller ion chains through photonic channels, to distribute computation across separate zones while maintaining coherence.

Recent work has shown that transport-based architectures, moving ions between storage and interaction zones, can partially mitigate mode crowding. However, these approaches introduce their own timing constraints and heating dynamics, and optimizing these trade-offs is now central to scaling beyond chains of 50–100 ions.

Another practical hurdle is the need for extensive laser infrastructure. Each ion requires carefully aligned beams for cooling, manipulation, and readout. This makes large-scale deployment technically demanding, although integrated optical systems, using on-chip waveguides and mirrors, are being actively investigated. Despite these constraints, the exceptional coherence of trapped-ion qubits (with some optical clock states remaining coherent for hours) and their high-precision gate operations have cemented this platform's place among the most promising candidates for early fault-tolerant computing.

In commercial systems, sustained performance is increasingly limited not by coherence but by calibration overhead. Small drifts in beam alignment or laser phase can lead to measurable shifts in multi-qubit gate phases, constraining long experimental runs and large-scale workloads.

Given their combination of accuracy, all-to-all connectivity, and long coherence, trapped-ion systems are strong contenders for implementing the first generations of error-corrected logical qubits.

C. PHOTONIC QUBITS

Photonic quantum computing utilizes light particles; typically, single photons to encode qubits through degrees of freedom such as polarization, path, or squeezed states. A notable benefit of this platform is that it can operate at room temperature, without the need for cryogenics, and is naturally suited for high-speed communication across optical fiber due to photons' low interaction with the environment and high transmission fidelity [13], [14]. Unlike matter-based systems, photonic qubits do not undergo decoherence in flight, maintaining their state unless absorbed or scattered. These properties have made photons the backbone of quantum networking and the basis of the first commercial quantum key distribution systems.

Although photonic qubits offer distinct strengths, building a general-purpose quantum processor with them presents unique hurdles. Many of these strengths are well established in controlled laboratory experiments, yet translating them into scalable, fault-tolerant systems remains difficult. Experimental demonstrations often rely on carefully optimized optical tables, whereas commercial devices must operate within integrated photonic platforms where losses, fabrication variability, and detector performance impose stricter practical limits. Since photons do not readily interact with each other, two-qubit operations must rely on indirect techniques; most commonly, entanglement generated through measurements or via coupling photons to intermediary systems such as atoms or quantum dots. Companies like PsiQuantum have adopted a measurement-based paradigm, where large-scale cluster states of entangled photons are generated and then measured to simulate gate operations [33]. This approach demands highly efficient photon sources, low-loss integrated optics, and ultra-fast detectors. PsiQuantum uses silicon photonic fabrication to produce chips with thousands of waveguides, aiming to construct photonic networks large enough for fault-tolerant computing. However, despite progress in silicon photonics, no photonic platform has yet demonstrated a fully error-corrected logical qubit, and existing prototypes operate at scales far smaller than those required for fault-tolerant computation.

Xanadu, on the other hand, has focused on continuous-variable photonics, utilizing squeezed light pulses. Its Borealis processor, featuring 216 modes, demonstrated quantum computational advantage in a specialized sampling problem in 2022 through Gaussian Boson Sampling; a task believed to be intractable for classical machines [34]. In this architecture, squeezed-state pulses pass through programmable interferometers arranged in optical fiber loops, producing interference patterns beyond classical reach.

Photonic system performance is typically evaluated via photon loss rates and interference visibility. Loss is the dominant source of error, while interference visibility directly impacts gate fidelity. Current lab-scale experiments report two-photon gate fidelities around 90–95%, with further improvements limited by detector inefficiencies and photon indistinguishability [32], [35]. Entanglement is often probabilistically generated (e.g., via spontaneous parametric down-conversion), which necessitates multiplexing many generation attempts and using fast feed-forward switching to maintain determinism. Furthermore, performance reported in short-duration, tightly controlled experiments often exceeds what can be sustained in long-running or cloud-accessible photonic systems. In contrast, long-running or cloud-accessible photonic devices must maintain stable source brightness, indistinguishability, and interferometer phase alignment over extended periods, which remains a major engineering obstacle.

Overcoming these limitations is key to scaling photonic quantum systems. It requires optical circuits that minimize losses during photon transmission and coupling; losses that can quickly add up across complex networks. Additionally, arranging photons to arrive at the right place and time with sufficient precision poses architectural challenges in synchronization and fast routing. However, ongoing research in integrated photonics, quantum-dot sources, and photonic error-correcting codes shows promise in overcoming these barriers.

Due to their immunity to thermal noise, intrinsic speed, and seamless compatibility with telecom infrastructure, photonic platforms are well-positioned for net-work-based quantum computing. Their strength in quantum communication and sampling tasks, like boson sampling and random circuit sampling, highlights their value as a core component in the broader quantum ecosystem.

D. NEUTRAL ATOMS

Neutral-atom quantum processors use individual atoms held in place by tightly focused laser beams, known as optical tweezers or arranged within optical lattices. Common atomic species include rubidium and cesium, chosen for their well-understood internal structure and optical transitions. Qubit states are typically encoded in two hyperfine levels of the atomic ground state. These energy levels are extremely stable, and in controlled environments, coherence times of several seconds have been demonstrated [25], [36].

A major strength of this platform lies in its scalability. Optical trapping techniques such as spatial light modulators and acousto-optic deflectors can generate large two-dimensional atomic arrays containing hundreds of qubits [36]. Being neutral, the atoms exhibit minimal interaction in their ground state, reducing unwanted coupling. However, by exciting selected atoms to Rydberg states, strong dipole–dipole or van der Waals interactions can be activated on demand, enabling entangling gates while pre-serving coherence in idle states [36].

Leading companies and academic groups have made significant progress. QuEra's 256-atom quantum simulator demonstrated programmable spin models, while Pasqal has explored quantum optimization algorithms, such as QAOA, implemented via the interaction graph of Rydberg atoms. Atom Computing's Phoenix system has shown over 1000 atoms loaded in a static array (though not all actively used in computation) [37]. This high qubit density, achieved without microfabricated circuits, highlights the architectural efficiency of neutral-atom platforms.

Gate-level control has also progressed: Rydberg-mediated two-qubit gates have reached fidelities around 97–98% in recent experimental demonstrations, with gate durations on the scale of a few microseconds [38]. Moreover, experiments have created GHZ states involving more than 20 atoms, showcasing the system's capacity for large-scale entanglement. Certain neutral-atom systems have implemented analog versions of QAOA, where the physical Rydberg interaction graph naturally encodes the optimization landscape [39].

Nonetheless, challenges remain. Improving two-qubit gate fidelity is essential for broader applications, particularly in digital quantum computing and error correction. Sources of error include atom temperature, laser phase noise, and field-induced level shifts. Mitigation techniques, such as enhanced cooling and tailored pulse sequences, are being actively researched [38].

Scalability is also limited by probabilistic trap loading: not all optical traps are reliably filled, with typical efficiencies between 50–90%. This necessitates post-selection or real-time rearrangement to achieve fully occupied arrays. Additionally, though atoms can be repositioned between runs, interactions are effectively local during execution, governed by the geometry and limited interaction range of Rydberg states ($\sim$10 μm). Strategies like dynamic atom movement allow for more complex connectivity but may introduce time overhead.

Readout is typically performed via fluorescence detection at the end of a computation. While accurate, this process is relatively slow (tens of milliseconds), and mid-circuit measurement remains an engineering challenge; though techniques such as shelving states offer possible solutions.

In larger neutral-atom arrays, long experimental runs can still be sensitive to slow drifts in trap depth, beam alignment and Rydberg-laser phase. These effects are routinely corrected through recalibration, but the need for frequent adjustments becomes more noticeable as systems grow in size. Maintaining uniform trapping conditions across many sites is technically demanding, and this difference becomes more apparent as systems move from short, controlled experiments to devices that must run reliably for hours or days.

Altogether, neutral-atom systems represent a promising route to scalable quantum processors with high qubit counts and long coherence. Their hybrid digital-analog capabilities, combined with flexible geometry and reprogrammable interactions, make them particularly attractive for quantum simulations and, with ongoing improvements, for universal fault-tolerant computing [36], [39].

E. SILICON SPIN QUBITS

Silicon spin qubits store quantum information in the spin orientation of individual electrons or nuclei, which are confined either in quantum dots or near donor atoms embedded in silicon. A key advantage of this platform is its close alignment with the existing semiconductor industry: fabrication techniques developed for classical integrated circuits, particularly CMOS processes, are directly applicable. This manufacturing synergy presents a credible path toward scaling up to extremely large qubit counts using tools already refined for producing advanced silicon chips [40].

A key material advantage of silicon is its availability in isotopically enriched form, particularly ^28Si, which has zero nuclear spin. When used as a host material, this allows for exceptionally long coherence times of spin qubits, with single electron spin lifetimes measured in the range of seconds under cryogenic and low-noise conditions [40]. Spin qubits are generally realized in two ways: gate-defined quantum dots; where single electrons are electrostatically confined using metallic gates on heterostructures (such as Si/SiGe), and donor qubits; where electrons or nuclei bound to dopant atoms (like phosphorus) serve as the qubit.

Over the past decade, several milestones have been achieved that mark the maturity of silicon spin qubit systems. These include high-fidelity single-shot readout of spin states, coherent control using electron spin resonance (ESR) or electric dipole spin resonance (EDSR), and the realization of two-qubit gates based on exchange interaction. By the end of the 2010s, experimental implementations of silicon-based spin qubits had reached high levels of control: two-qubit gate fidelities exceeding 98% were first demonstrated in 2019, with further improvements pushing fidelities to around 99.5% by 2022. These values are comparable to the best results achieved in other leading platforms such as superconducting circuits and trapped ions [40].

In parallel, Intel applied its large-scale fabrication capabilities to create the Tunnel Falls chip, a 12-qubit silicon device built using standard 300 mm CMOS processing lines, a step intended to test reproducibility and uniformity of spin qubit arrays [41]. This prototype aims to validate high-yield, uniform silicon qubit production and test automated control strategies for scaling up [41]. At the same time, academic teams at Delft University of Technology, UNSW Sydney, and CEA-Leti have demonstrated 4–6 qubit arrays capable of running basic quantum algorithms such as Grover search and quantum error detection protocols.

A major strength of silicon spin qubits lies in their exceptionally small footprint: individual quantum dots can span only a few tens of nanometers, allowing many qubits to be integrated within a limited area. Their control relies on fast electrical or magnetic pulses, making it possible to perform gate operations on nanosecond time-scales. This combination of size and speed positions spin qubits as a leading candidate for high-density, low-latency quantum processors. Silicon's ubiquity in classical computing also raises the possibility of co-integration of quantum and classical logic on a single chip; a feature that could help reduce interconnect overhead and facilitate on-chip feedback and control [40].

An exciting development has been the use of superconducting microwave resonators to mediate interactions between distant spin qubits separated by more than 10 μm. This approach circumvents the limitations of nearest-neighbor exchange coupling and provides a scalable path to modular quantum processor architectures [42].

In recent multi-qubit demonstrations, device performance has increasingly been limited not by coherence but by variability between nominally identical quantum dots. Small shifts in electrostatic environment can alter resonance frequencies by several megahertz over the course of a single experimental session, requiring periodic retuning. These drifts do not usually prevent operation but place noticeable demands on calibration procedures, especially when operating larger arrays. In addition, reported fidelities in laboratory settings are often obtained under tightly optimized tuning conditions; maintaining similar performance on continuously operating or remotely accessed devices remains an open engineering challenge.

However, there are significant technical challenges. Unlike ions or atoms, no two silicon quantum dots are identical. Variations in fabrication lead to differences in confinement potentials, tunnel barriers, and qubit frequencies. As a result, tuning and calibrating multi-qubit arrays is labor-intensive and sensitive to charge noise and material imperfections. While advanced device engineering and sweet-spot biasing help reduce noise-induced dephasing, the reproducibility and uniformity of qubit behavior remain open concerns for large-scale deployment.

Two-qubit operations in silicon often rely on fine-tuning the exchange interaction, which is highly sensitive to fluctuations in gate voltages and cross-talk between neighboring control electrodes. This limits operation robustness and places stringent demands on gate design and filtering. Readout, too, can be a bottleneck: spin qubits typically require nearby charge sensors or resonators, and multiplexing large arrays of readout channels remains a challenge for system scalability.

Despite these obstacles, the long-term appeal of silicon spin qubits remains strong. If quantum error correction protocols demand hundreds or thousands of physical qubits per logical qubit, then the small footprint and fast control offered by silicon may prove decisive. Furthermore, if control electronics and quantum circuits can be monolithically integrated, silicon might offer one of the few viable paths to building truly large-scale, fault-tolerant quantum computers. Preliminary results from Intel's Tunnel Falls project suggest that industrial fabrication can yield highly uniform and reproducible qubit arrays [41], [43].

F. TOPOLOGICAL QUBITS

Topological qubits represent one of the most conceptually ambitious directions in quantum computing. Their design stems from topological phases of matter; exotic quantum states where information is stored globally rather than locally. The leading physical candidate for realizing such qubits is the Majorana zero mode, a quasiparticle that is its own antiparticle. Majorana modes are predicted to emerge under specific conditions in hybrid nanostructures, such as semiconductors with strong spin–orbit coupling (e.g., InAs or InSb) that are proximity-coupled to conventional superconductors [44], [45].

What makes Majorana-based qubits theoretically compelling is the topological protection they promise. A qubit state can be encoded nonlocally in a pair of spatially separated Majoranas. Because local noise typically cannot affect both ends of the qubit simultaneously, such an encoding is inherently resistant to many common sources of decoherence. Operations on these qubits would proceed via braiding; moving the Majoranas around one another in space, which alters the qubit state in a way that depends only on the topological class of the braid, not the details of the physical motion [44], [45].

Among industrial players, Microsoft has taken a leading role in attempting to realize Majorana-based qubits. Their experimental program has focused on fabricating epitaxial nanowires, typically indium arsenide or antimonide wires covered with superconducting aluminum layers, and tuning them to exhibit signatures of Majorana modes. In 2022–2023, researchers affiliated with Microsoft reported satisfying several criteria believed to be consistent with a topological superconducting phase, including observations aligned with the topological gap protocol [46]. However, these results remain the subject of intense scrutiny, and no experimental demonstration of qubit-level functionality (e.g., initialization, control, readout, or braiding) has yet been achieved.

Should topological qubits prove experimentally viable, they may offer a transformative path for minimizing the scale of quantum error correction. In conventional qubit architectures, achieving fault tolerance typically involves layering extensive software-level correction over physical qubits; often requiring the coordination of hundreds or even thousands of them to encode a single logical qubit. By contrast, the distinctive structure of topological qubits embeds error resilience at the hardware level: quantum information is encoded nonlocally across spatially separated quasiparticles, making it inherently less vulnerable to many types of local disturbances. This intrinsic protection has the potential to significantly reduce the number of required physical components in future large-scale quantum systems. Theorists have suggested that a logical qubit could be constructed from just 10–20 Majorana modes, dramatically shrinking the footprint of scalable quantum architectures [45].

Nonetheless, the challenges are formidable. Demonstrating true Majorana modes, as distinct from trivial Andreev bound states or other zero-energy artifacts, is itself a major scientific hurdle. Proving non-Abelian statistics via controlled braiding is even more difficult and requires highly sophisticated nanowire junctions, precise control over gate voltages and magnetic fields, and ultra-low temperature operation [44]. Furthermore, current proposals for universal quantum computation often require non-topological operations to supplement braiding, such as magic-state distillation, which introduces additional complexity [45].

Microsoft has adopted a diversified approach, building a software ecosystem (Azure Quantum) to interface with other quantum hardware platforms while continuing their long-term push toward Majorana-based systems. As it stands, topological qubits remain a high-risk, high-reward bet. Their theoretical benefits are unmatched in terms of hardware-level fault tolerance, but the experimental pathway is still uncertain and under active investigation.

Although the theoretical framework behind Majorana-based qubits is well established, experimental signatures reported so far remain sensitive to device design, material interfaces, and data-analysis methods. Several studies in the broader community have shown that zero-bias peaks and related spectroscopic features can arise from non-topological mechanisms, underscoring the importance of reproducible measurements across different fabrication runs. Moreover, most demonstrations to date involve static spectroscopy rather than active qubit manipulation, and scaling such devices into multi-mode networks introduces further constraints on junction uniformity, magnetic-field stability, and quasiparticle poisoning. These practical factors presently limit the rate at which new demonstrations can be independently verified.

Despite these technical hurdles, the pursuit of topological qubits remains one of the most ambitious and conceptually transformative directions in quantum computing. Their potential to provide native error resilience continues to drive experimental progress, even if functional qubits remain elusive. To situate this approach within the broader spectrum of available quantum hardware platforms, Figure 2 provides a comparative visual landscape. The comparison is based on representative ranges commonly reported in the literature and is intended to reflect reported best-case or milestone results, rather than sustained average performance. Accordingly, gate fidelities refer to single- and two-qubit gate fidelities under optimized experimental conditions, while descriptors such as ‘room temperature’ for photonic platforms are used to highlight their operational environment rather than to specify a characteristic coherence-time scale.

As each qubit platform brings its own trade-offs in terms of control, fidelity, and scalability, it becomes essential to develop consistent benchmarks to evaluate their capabilities. This naturally invites closer scrutiny of how performance across architectures is measured and compared.

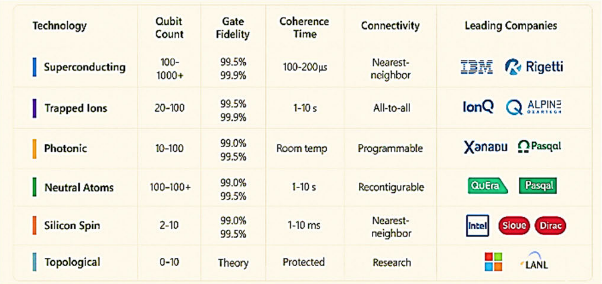

Technology	Qubit Count	Gate Fidelity	Coherence Time	Connectivity	Leading Companies
Superconducting	100- 1000+	99.5% 99.9%	100-200µs	Nearest-neighbor	IBM Rigetti
Trapped Ions	20-100	99.5% 99.9%	1-10 s	All-to-all	IonQ ALPINE
Photonic	10-100	99.0% 99.5%	Room temp	Programmable	Xanadu Pasqal
Neutral Atoms	100-100+	99.0% 99.5%	1-10 s	Recontigurable	QuEra Pasqal
Silicon Spin	2-10	99.0% 99.5%	1-10 ms	Nearest-neighbor	Intel Sloue Dirac
Topological	0-10	Theory	Protected	Research	LANL

FIGURE 2. Comparative landscape of major quantum hardware platforms, generated using claude 3.5 (Anthropic, 2025).

III. KEY METRICS FOR EVALUATING QUANTUM PROCESSORS

Evaluating the capabilities of quantum processors involves more than just tallying qubits. The hardware characteristics and system-level differences outlined in Section II and summarized in Table 1 motivate the need for consistent and architecture-aware benchmarking methodologies. Although qubit count is frequently highlighted in public discourse, it often offers a limited view of a processor’s true potential. A more robust assessment must include other critical factors like gate and readout fidelity, error rates, coherence durations, operation speeds, and inter-qubit connectivity. Taken together, these parameters offer a more realistic indication of how well a QPU can perform meaningful quantum computations.

At present, each quantum hardware platform demonstrates strengths in certain areas while facing limitations in others. Superconducting qubits, for example, lead in terms of sheer scale and gate speed, whereas trapped ions demonstrate exceptional coherence and fidelity. Meanwhile, emerging architectures like neutral atoms and photonic systems are closing the gap by advancing in error correction and system integration. This technological diversity creates the need for standardized benchmarks that can allow apples-to-apples comparisons between disparate systems. Metrics such as Quantum Volume (QV) and Algorithmic Qubits (#AQ) have thus emerged as integrative tools to reflect the combined effect of various physical parameters. Many current benchmarking efforts are therefore increasingly oriented toward demonstrating quantum utility rather than isolated advantage experiments.

Although these metrics capture broad trends, they can differ noticeably between short, optimized calibration runs and long-duration operation. Reported peak fidelities and coherence times often reflect conditions tuned for a specific experiment, whereas cloud-deployed or continuously running devices must maintain performance over many hours or days. As a result, nominal values in specification sheets can overestimate the performance achievable on arbitrary workloads, making temporal stability an increasingly important part of evaluating processor capability.

Operating conditions play a significant role in shaping the design and practicality of quantum systems. For instance, superconducting-based architectures typically demand dilution refrigerators to maintain millikelvin temperatures, whereas photonic platforms can function at room temperature. These operational requirements affect not just technical performance but also influence engineering complexity, energy consumption, and system maintainability. Moreover, the pace of technological improvement varies across modalities: superconducting systems benefit from existing CMOS infrastructure, whereas topological or photonic systems often require novel fabrication pathways and materials science advances. Therefore, tracking progress over time becomes just as important as static benchmarking.

In practice, the rate at which these metrics improve is not uniform across platforms. Superconducting systems often exhibit steady, incremental gains driven by fabrication and control improvements, whereas neutral-atom and photonic architectures may show more abrupt progress linked to architectural innovations. A crucial factor in these assessments is scalability, defined as the ability to increase qubit numbers while maintaining performance and stability. However, larger systems introduce additional calibration complexity, crosstalk, and parameter variability, which contribute to the widening gap between peak and sustained performance.

As systems increase in size, cross-talk, calibration overhead, and parameter drift typically grow as well. Larger devices often require more frequent recalibration cycles, and inter-qubit variability becomes more visible, especially for architectures relying on analog control. These effects do not prevent operation but contribute to a widening gap between peak and average performance, which is an important consideration when assessing scalability claims.

Regular calibration and diagnostic testing are essential for maintaining consistent performance in quantum devices. Techniques like randomized benchmarking, cross-entropy benchmarking, and process tomography are commonly used to measure gate accuracy and operational reliability under realistic usage. On top of that, innovations in software, such as the use of logical qubits or compiler-driven virtual mappings, aim to compensate for physical limitations, introducing additional layers of complexity in system evaluation.

Benchmarking results must also be interpreted in light of the specific calibration state of the device. Some metrics, such as randomized benchmarking, are relatively insensitive to coherent errors but may overlook crosstalk or drift over time. Others, such as tomography or cycle benchmarking, provide more detailed information but are expensive to run at scale. Consequently, a single benchmark rarely captures the full operational profile of a QPU.

Beyond hardware characteristics, a processor's practical performance often depends on its software ecosystem. The quality of the compiler, support for error mitigation, and integration with development frameworks can all enhance real-world out-comes. As a result, even devices with slightly lower fidelity or connectivity may deliver better algorithmic performance if their control stack efficiently adapts to hardware limitations. Consequently, system-level integration, across hardware, firmware, and software emerges as a significant contributor to end-to-end quantum performance.

In multi-user environments, compiler behavior and queueing conditions can meaningfully influence observed performance. Devices with strong software stacks may compensate for hardware non-uniformity through optimized routing and error mitigation, whereas less mature stacks can underutilize otherwise capable hardware. These system-level interactions increasingly shape practical performance outcomes and complicate direct hardware-only comparisons.

Furthermore, reproducibility and stability over time are critical. Metrics gathered during initial calibration runs can vary significantly with environmental drift, hardware aging, or recalibration cycles. Thus, benchmarks must often be averaged over multiple runs or reported alongside uncertainty measures. Temporal consistency in achieving benchmark scores like Quantum Volume or Algorithmic Qubits indicates not just performance, but also engineering maturity and reliability.

Long-term reproducibility remains an active challenge for all platforms. Sequential runs of the same benchmark may yield slightly different outcomes due to calibration drift, thermal fluctuations, or control-electronics noise. For this reason, many research groups now report performance as a range or as a distribution over multiple trials, providing a more realistic picture of day-to-day device behavior.

Ultimately, the goal is to assess how far a given QPU has moved toward achieving practical quantum advantage; the point at which it can outperform classical computers on relevant tasks. To this end, both empirical and algorithm-specific benchmarks are being increasingly emphasized by developers and researchers alike. These include running quantum chemistry simulations, solving optimization problems, or validating output distributions in random circuit sampling.

Below are the core performance metrics typically used in comparative evaluations:

- Qubit Count: The most basic metric is the number of physical qubits. This number is often highlighted in press releases (e.g., IBM's 127-qubit Eagle or IonQ's 32-ion systems) [10], [17]. However, qubit count alone can be misleading; 100 noisy qubits are not necessarily better than 10 high-quality qubits. Still, qubit count sets an upper bound on the size of quantum circuits (in terms of width) that can be executed. As seen in Table 1, current QPUs range from under 100 qubits (ion traps, spin qubits) to several hundred (superconducting, photonic modes) as of 2023, with some platforms targeting 1000+ in the near future.
- Qubit Quality (Fidelity and Coherence): Arguably more important than sheer quantity is the quality of qubits. Gate fidelity refers to the accuracy of quantum gate operations. It is typically measured by randomized benchmarking or cycle benchmarking, yielding a

probability of error per gate. High fidelity (close to 100%) is essential for deep circuits. As of 2023, state-of-the-art one-qubit gate fidelities are >99.9% for many platforms (superconducting, ions, spin), and two-qubit gate fidelities have exceeded 99% in several systems [10], [27], [29]. Readout fidelity (the accuracy of measuring a qubit's state) is another aspect, often slightly lower than gate fidelity (e.g., 95–99% typical). Coherence time is the time over which a qubit maintains its quantum state. There are two relevant times: T_1 (relaxation time) and T_2 (dephasing time). A useful rule of thumb is that a qubit should be able to undergo many operations within its T_2 time. For example, if $T_2 = 100$ μs and gate time is 20 ns, then about 5,000 operations can be performed before dephasing; enough for certain algorithms but not for full error correction. Trapped ions have extremely long coherence (seconds to minutes) [25], whereas superconducting qubits have shorter T_1, T_2 (tens to hundreds of μs) and thus require faster gates to compensate [40], [43].

- Connectivity: This describes which qubits can interact with each other directly. Some architectures, like trapped ions or certain photonic schemes, have effectively all-to-all connectivity (any qubit can be entangled with any other via a single operation) [29]. Others, like superconducting qubits on a chip or neutral atoms in a grid, have nearest-neighbor (local) connectivity; qubits only directly interact with adjacent qubits, and longer-range interactions require intermediate swaps. Limited connectivity can significantly increase the number of gates required to implement a given algorithm, as extra SWAP gates must be inserted to move quantum information around, each introducing more error. Therefore, improving connectivity (either physically or through smart compilation) is a key engineering focus. Some superconducting designs use crossovers or 3D wiring to go beyond 2D nearest-neighbor coupling. In neutral atoms, dynamic reconfiguration of the array can effectively make the interaction graph more flexible.
- Gate Speed: The duration of qubit operations is another critical metric. Faster gates mean more operations can be completed before decoherence sets in. Super-conducting qubits currently lead in speed, with gate times on the order of 10–100 ns. Spin qubits are also fast (tens of nanoseconds for single-qubit gates, a bit longer for two-qubit). Photonic gates (in measurement-based schemes) can be nano-second-scale as well for optical circuit transformations. Trapped ions are slower: single-qubit gates ~1 μs, two-qubit gates 10–100 μs. Neutral atoms with Rydberg gates are intermediate (~0.1–1 μs). While not as prominently reported as fidelity or qubit count, speed is part of the "quantum compute power" of a device. A processor with slow gates might require more careful error mitigation since qubits can decohere during long operations. Speed also affects clock cycle time for algorithms. That said, extremely fast gates can sometimes induce more crosstalk or error (e.g., faster pulses may cause leakage out of the qubit state in transmons), so each platform balances speed and fidelity.
- Error Rates: Beyond individual gate fidelities, one can speak of error rates per operation or per circuit. This includes depolarizing error rates (random errors in gates), readout error rates, and leakage rates (probability of a qubit leaving the computational subspace). Typical two-qubit gate error rates in 2023 are in the 10^{-3} to 10^{-2} range (0.1%–1%) for many systems [43], though some ion traps and super-conducting gates have dipped below 10^{-3}. Error rates need to be pushed down to the 10^{-4} or 10^{-5} range (or lower) for the surface code and other quantum error correction schemes to work effectively without unfeasible overhead [22], [42].

To address the complex interdependencies among qubit performance parameters, composite benchmarking approaches have been introduced:

- Quantum Volume (QV): Developed by IBM, this metric assesses a system's capability by determining the largest square-shaped random quantum circuit it can execute reliably, where the number of layers equals the number of qubits involved [17]. QV does not assess each hardware feature in isolation; instead, it captures how different aspects, such as the number of qubits, quality of gate operations, and the level of qubit interconnectivity, work together to determine overall device performance. A system with higher QV is generally capable of executing deeper and wider quantum circuits while keeping error rates manageable, thereby providing a more realistic indicator of computational capability. Each increase in QV often requires improving multiple aspects of the device. For example, IBM's 27-qubit Falcon r5 achieved QV 32 in 2019, and by 2021 a 127-qubit system (Eagle) reached QV 128 by both adding qubits and keeping error rates low. Quantinuum's 20-qubit H1-1 achieved a record QV 4096 in 2021 by leveraging very high fidelities [19].
- Algorithmic Qubits (#AQ): Proposed by IonQ, #AQ is defined via performance on a specific suite of algorithms, giving a sense of how many usable qubits the system has for real tasks [10]. In practice, #AQ is determined by running algorithms (like variational optimization, Monte Carlo, Fourier transforms, etc. as per IonQ's description) and seeing how many qubits can be involved before results deteriorate below a threshold [10]. IonQ reported #AQ 25 for Aria and #AQ 29 for Forte, meaning those machines could run certain reference algorithms on 25 or 29 qubits with >37% success probability [10].
- Quantum Advantage Benchmarks: Aside from general metrics, specific demonstrations of quantum advantage (solving a problem faster or better than a classical computer) serve as important benchmarks. Examples include random circuit sampling (Google's supremacy test), boson sampling (Xanadu's Borealis), or solving

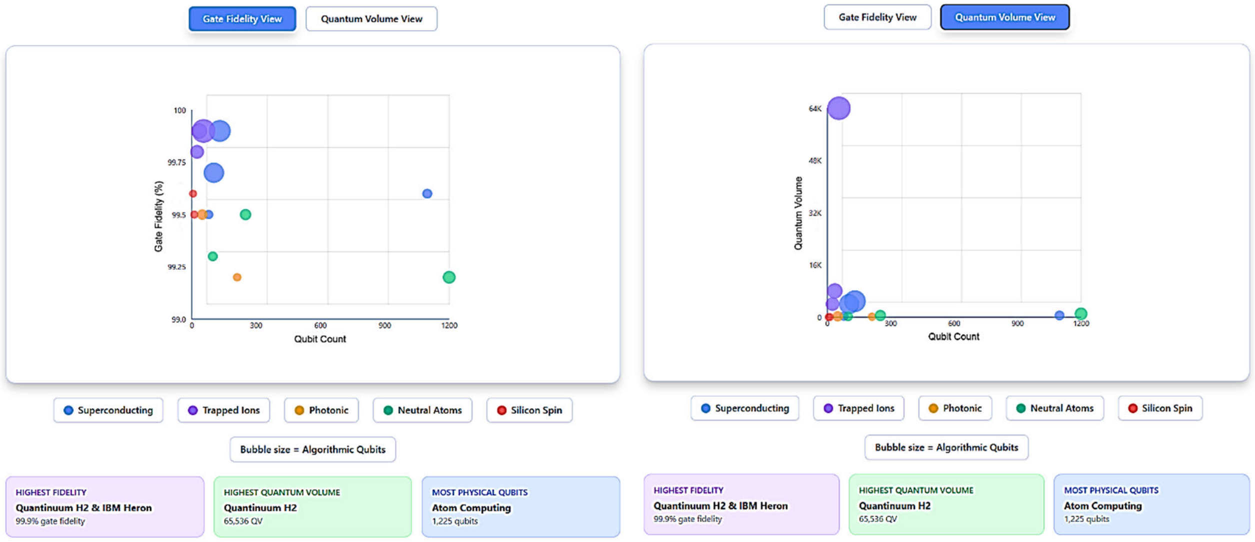

FIGURE 3. Dual-metric quantum processor performance map (gate fidelity and quantum volume views), generated using claude 3.5(Anthropic, 2025).

specific algebraic problems (e.g., factoring 153 on a trapped-ion system) [5], [9], [34]. These results signal technological maturity for specific workloads and illustrate the multifaceted nature of QPU development.

To better understand how different platforms balance core hardware trade-offs, such as qubit number and gate fidelity, a comparative visual summary is shown below.

Superconducting platforms, such as IBM's Condor and Willow, demonstrate scalability with qubit counts in the hundreds, though often at the cost of slightly reduced gate fidelities (∼99–99.5%). In contrast, trapped-ion systems maintain near-ideal fidelities (approaching 99.9%) but currently operate with more modest qubit counts. Neutral-atom devices show intermediate performance, balancing fidelity and scale, while photonic platforms support high qubit numbers yet continue to lag behind in gate quality. The large representation of Algorithmic Qubits highlights platforms that offer strong real-world performance through effective integration of hardware and software layers.

The distribution of platforms illustrates the trade-offs each architecture faces; between scale, precision, and connectivity. The diagonal trajectory across the map indicates the industry's collective push toward systems capable of supporting both large-scale quantum circuits and high-fidelity operations. Reaching the upper-right corner remains a defining goal in the path toward fault-tolerant and application-ready quantum computing.

While current metrics help characterize the progress made across various quantum platforms, they also bring into focus the technical depth and engineering nuance required to move forward. Rather than offering definitive conclusions, these indicators point to areas where further refinement, stability, and innovation will be essential. The following chapter explores such aspects in more detail, focusing on the ongoing technical complexities facing quantum computing research and development.

IV. CURRENT CHALLENGES AND FUTURE PERSPECTIVES

Despite significant technological progress over the past decade, current quantum processors still operate within the regime characterized by John Preskill as NISQ – Noisy Intermediate-Scale Quantum computers [4]. The platform-specific strengths and limitations discussed in Sections II and III directly shape the engineering and scalability challenges examined in this section. This term encapsulates two key limitations: ''Noisy'' refers to the inherently fragile nature of quantum bits, which suffer from decoherence and operational errors; ''Intermediate-Scale'' indicates that today's machines consist of tens to several hundred qubits; beyond the reach of brute-force classical simulation for arbitrary quantum states, but far from the scale required for large-scale, fault-tolerant quantum computing.

Several of the following challenges are closely interconnected and often manifest differently across hardware platforms. Therefore, recurring themes such as noise, scalability, and calibration stability are revisited from complementary engineering and system-level perspectives.

In practice, current quantum hardware is constrained by limited coherence, imperfect gate operations, and noisy readout, which restrict circuit depth and algorithmic complexity. While proof-of-concept demonstrations have been achieved, large-scale applications remain dependent on further improvements in reliability and error suppression [5].

These limitations underscore the essential distinction between quantum feasibility and quantum utility. While feasibility has been established, showing that quantum systems can outperform classical ones in tightly defined scenarios, utility requires the capacity to execute error-resilient, application-relevant algorithms on programmable hardware.

This distinction reflects the broader shift in the field from isolated demonstrations of advantage toward sustained and application-oriented quantum utility.

Moreover, the landscape of quantum technologies is increasingly heterogeneous. Each hardware platform: superconducting circuits, trapped ions, photonic processors, neutral atoms, spin qubits, exhibits distinct performance characteristics and engineering trade-offs [13]. For instance, superconducting devices have scaled fastest in terms of qubit number, while ion traps lead in gate fidelity and coherence times. While individual platforms have demonstrated strong performance in specific areas, such as gate fidelity, qubit coherence, or system size, none currently delivers all these attributes in unison. Variability in platform performance makes it difficult to draw meaningful comparisons between technologies. Benchmarking tools must therefore account for differing trade-offs in fidelity, connectivity, and circuit depth, which often limits their universality.

Reaching the stage of fault-tolerant quantum computing will require advances on multiple levels; prolonging coherence through improvements in device fabrication, refining interfaces between quantum hardware and classical control logic, and deploying error correction protocols that can be realistically scaled with system size. Advancing quantum computing is not solely a matter of adding more qubits; it demands a coordinated engineering approach that integrates improvements in physical hardware, control software, and algorithmic strategies.

Building scalable and reliable quantum processors involves a multi-pronged effort; addressing noise and decoherence, refining hardware architecture for large-scale integration, and advancing techniques for active error correction to ensure stable computation over long runtimes. Alongside these efforts, it is equally important to explore promising application areas that could deliver practical value from near-term quantum systems, even before full fault tolerance is achieved. The following discussion presents a structured overview of these intertwined challenges and the directions in which current research is moving.

A. DECOHERENCE AND ERROR MITIGATION

Having previously introduced decoherence and gate fidelity as core performance considerations for quantum processors, we now delve more deeply into their physical origins, engineering implications, and the strategies emerging to mitigate their impact. Decoherence, the gradual loss of quantum information due to environmental interactions, remains the central obstacle to scalable quantum computation. Even weak coupling to external degrees of freedom, such as thermal fluctuations or electromagnetic noise, can lead to phase or energy relaxation, thereby limiting algorithmic depth and reliability.

Different physical implementations exhibit distinct sources of decoherence. For example:

- Superconducting qubits are vulnerable to material defects, particularly two-level system (TLS) impurities at interfaces and dielectric loss. Electromagnetic noise, crosstalk, and fluctuations in bias currents further degrade coherence. Advances in fabrication, such as cleaner substrates, improved lithography, and elimination of lossy interfaces, have extended coherence times, but performance remains limited to the sub-millisecond regime [6], [40].
- Trapped ions exhibit exceptionally long coherence times, on the order of seconds or more, owing to their isolation in ultra-high vacuum. However, they are still susceptible to decoherence via magnetic field drifts and motional heating. Engineering solutions such as "clock" transitions (magnetically insensitive states) and symmetric trap designs have proven effective in mitigating these effects [9], [18].
- Spin qubits, including those based on silicon or diamond NV centers, can decohere through charge noise or hyperfine interactions with surrounding nuclear spins. Isotopic purification (e.g., using ^{28}Si) significantly reduces magnetic noise, pushing coherence times into the millisecond range [24], [28], [40].
- Neutral atoms and photonic platforms face other limitations. Laser phase noise and black-body radiation can disrupt atomic states, while photon loss and mode mismatch challenge optical systems. These effects are being addressed through cavity-QED techniques, cryogenic cooling, and fault-tolerant photonic encoding [13], [14], [34].

Despite continued material and design improvements, today's gate operations still exhibit average error rates in the range of 0.1% to 1% (10^{-3} to 10^{-2}) per gate. These rates, while impressive from a physics standpoint, remain too high for reliable execution of deep quantum circuits. For example, a quantum algorithm requiring 1,000 gate layers would likely fail with near certainty if each layer introduces a 0.5% error. Thus, mitigation strategies, short of full error correction, have become an essential research focus [16], [25].

A variety of error mitigation techniques have emerged for current Noisy Intermediate-Scale Quantum (NISQ) devices:

- Zero-noise extrapolation artificially amplifies error rates and fits the outcomes to estimate the zero-noise limit. This technique assumes the noise is well-characterized and behaves linearly over the extrapolation range.
- Probabilistic error cancellation employs knowledge of the noise model and measurement outcomes to construct classically weighted ensembles that "cancel out" known error channels, albeit at the cost of exponentially increasing sampling requirements.
- Symmetry verification and post-selection strategies exploit conserved quantities in quantum simulations (e.g., total particle number or spin) to filter out unphysical outcomes.

While these approaches have shown promise in enhancing results for variational quantum eigensolvers (VQE) and other hybrid algorithms, they are inherently limited to shallow circuits, as overheads increase rapidly with circuit depth and qubit count [16].

Crucially, recent breakthroughs are beginning to signal a shift from mitigation toward true error correction, where redundancy is used to encode logical qubits capable of sustaining information over many physical cycles. In 2023, Google demonstrated that increasing the code distance in a surface code, from 3 to 5, led to a measurable drop in logical error rates. Specifically, encoding a logical qubit using 49 physical qubits (distance-5) outperformed a 17-qubit (distance-3) code, crossing the long-theorized fault-tolerance threshold around 99% fidelity per gate [23], [28].

Similarly, IBM achieved a significant milestone using its 127-qubit Eagle device: a logical qubit implemented with a heavy-hexagon code exhibited longer lifetimes than any of the underlying physical qubits [47], [48]. This result, though still preliminary, demonstrates the first empirical step toward "error suppression by scaling," where adding physical qubits improves rather than degrades overall reliability.

These advances represent more than technical victories; they serve as proof-of-principle that fault-tolerant quantum computation is not merely theoretical. Still, achieving the low error rates required for practical quantum error correction remains out of reach for most current systems. Encoding a single robust logical qubit may demand thousands of physical qubits, depending on the architecture and the code used, posing a major scalability bottleneck [49].

Although these demonstrations mark an important milestone, they remain limited in scope. Logical-qubit lifetimes are typically characterized under carefully chosen circuits and relatively short sequences, and results are often reported for a small number of logical qubits encoded in a single device. At present, there is still limited published data on how such schemes perform under sustained algorithmic workloads or across multiple hardware generations, so extrapolating these early results to full-scale fault-tolerant architectures should be done with caution.

Rather than depending on a singular breakthrough, advances in quantum reliability are likely to stem from the steady accumulation of improvements across many fronts. Enhancements in device fabrication, refinement of noise control techniques, and the use of tailored mitigation methods are all contributing incrementally. These technical gains, when combined with progress in calibration protocols, hardware control systems, and software that adapts to device-specific behavior, are gradually paving the way toward more stable and usable quantum systems.

B. SCALABILITY AND ENGINEERING CHALLENGES

Scaling quantum computers from tens to thousands of qubits introduces a diverse set of engineering hurdles, many of which have no direct analog in classical computing. One of the most immediate obstacles is wiring and control complexity. In current superconducting and semiconductor spin-qubit systems, each qubit typically requires dedicated microwave and DC lines for control and readout. As qubit numbers increase, this requirement becomes physically unsustainable due to space constraints and thermal loading within the dilution refrigerator environment. To address this, research is actively exploring multiplexing strategies such as frequency-division multiplexing and time-domain sharing, along with the development of cryogenic classical control electronics that can be colocated with the qubit chips inside the cryostat to reduce latency and heat load [6], [7], [41].

Ion trap and neutral atom platforms face a different scaling bottleneck: delivering and stabilizing hundreds of laser beams for individual qubit addressing. Proposed solutions include integrated photonic circuits for beam routing and magnetic field gradients to enable frequency-selective addressing. These technologies aim to reduce system footprint and enhance parallelism, though challenges in maintaining optical phase coherence and alignment remain [9], [10], [26].

Thermal management is another critical concern. As more control hardware is packed near the qubit substrate, heat dissipation from both classical electronics and quantum processes becomes non-negligible. Maintaining the sub-20 millikelvin environment necessary for coherent qubit operation requires highly efficient cryogenic infrastructure and materials with low thermal conductivity yet high electromagnetic compatibility [6], [17].

Fabrication yield and qubit uniformity become increasingly important as chip sizes grow. Even minor process variations in junction resistance, oxide thickness, or lithographic patterning can cause significant disparities in qubit frequency and performance. In silicon quantum dot systems, for instance, variability in gate voltage thresholds introduces complex tuning requirements that do not scale well. Automated calibration and machine learning-based tuning protocols are under investigation to address this issue [40], [41]. In superconducting circuits, frequency crowding and cross-talk between resonators are growing concerns, especially for large monolithic chips such as IBM's planned 1121-qubit Condor processor, which uses 3D integration and chiplet tiling to manage density and interconnectivity [7].

In many of these scaling proposals, the underlying assumptions about calibration overhead and device-to-device variability are still being tested. Early multi-chip or multi-module experiments often reveal additional sources of error at interfaces, such as mode mismatch, packaging-induced strain, or timing skew between control channels. These effects are seldom captured in small-scale prototypes but can become prominent once systems move beyond a few tens or hundreds of qubits.

A promising approach to overcoming physical scaling limits involves modular quantum architectures. These systems comprise smaller, manageable quantum processing units (QPUs) that are interconnected to form a larger machine. Interconnects may be realized via optical photons (suitable for long-range entanglement, especially in ion or atom systems) or via microwave resonators (more appropriate for superconducting circuits within a shared cryostat). While optical links offer scalability, photon loss and coupling efficiency remain limiting factors; current rates of inter-module entanglement are on the order of a few Hertz [10]. Conversely, microwave links promise higher rates but are generally constrained to short distances [6], [42].

Scalability also depends on how well software adapts to increasingly complex hardware. As quantum processors grow, circuit compilation becomes more constrained by qubit connectivity, coherence times, and native gate sets. Techniques like qubit mapping, adaptive gate synthesis, and real-time error-aware compilation are being developed to better align quantum programs with underlying hardware characteristics [35]. Simultaneously, classical simulation of quantum systems for verification and optimization becomes computationally intensive. Hybrid methods using tensor networks, quantum-inspired algorithms, and approximation schemes may offer feasible paths forward for debugging and benchmarking near-term systems [1], [8].

In summary, scaling quantum systems is a multi-dimensional problem involving not just qubit number but also interconnect design, thermal engineering, fabrication reliability, and software adaptability. Addressing these interdependent challenges will require coordinated advances in quantum hardware design, materials science, control infrastructure, and software ecosystems. The experience from classical high-performance computing offers useful guidance, but entirely new paradigms are likely needed to reach fault-tolerant and application-relevant quantum computation.

As a result, reported roadmaps often represent best-case engineering targets rather than guaranteed timelines. Iterative feedback between prototype hardware, control electronics and application-level testing typically determines which scaling approaches remain viable, and some currently popular architectures may need to be substantially revised as new bottlenecks appear at larger system sizes.

C. TOWARD FAULT TOLERANCE: ERROR CORRECTION ROADMAP

While previous sections outlined the physical and engineering constraints of current quantum devices, this section explores the structured progression toward fully fault-tolerant quantum computation.

Achieving fault-tolerant quantum computation remains the defining challenge in the pursuit of scalable quantum systems. Fault tolerance refers to the ability to execute arbitrarily long quantum algorithms while actively suppressing errors through the encoding of logical qubits. These logical units are formed from many physical qubits and protected via continuous error detection and correction cycles that preserve quantum information without measurement collapse [3], [23].

The theoretical underpinnings of quantum error correction (QEC) are mature. Surface codes and quantum Low-Density Parity-Check (LDPC) codes are leading candidates due to their high fault-tolerance thresholds, which range between one error in a thousand and one error in a hundred gate operations [23], [47]. Only recently have physical platforms approached these thresholds, setting the stage for early demonstrations of fault-tolerant primitives.

Initial progress between 2020 and 2022 featured Google's implementation of a distance-3 surface code on a superconducting platform and experiments using ion traps and photonic systems to correct individual errors [23], [28]. Initial demonstrations provided useful insights into the implementation of quantum error correction, although logical qubits generally exhibited higher error rates than physical ones during this phase. In 2023, experiments involving larger surface codes, such as the distance-5 configuration, indicated more effective error suppression compared to smaller implementations [28]. Around the same period, IBM demonstrated that logical qubit states could persist longer than their physical counterparts by leveraging repetition coding techniques [48]. As can be seen in Figure 4, the trajectory from early error correction demonstrations to logical qubit systems is marked by a sequence of increasingly demanding milestones.

These visualized transitions highlight how each step, from basic repetition codes to multi-logical qubit arrays, requires substantial architectural and algorithmic refinement before true fault tolerance can be achieved.

The focus from 2024 to 2026 is expected to include demonstrations of logical memory outperforming the best physical qubits and fault-tolerant logical gate operations, such as a logical CNOT, that exceed the fidelity of their physical counterparts [48], [49].

Rather than representing a single breakthrough, the roadmap toward fault-tolerant computation involves a sequence of increasingly sophisticated capabilities. Looking ahead to the latter half of the 2020s, a major objective involves building compact systems composed of multiple logical qubits that can interact coherently. These subsystems may support basic quantum algorithms or serve as error-protected units within larger architectures. Realizing such configurations, particularly with target logical error rates around one in a million, is expected to require substantial physical resources, potentially over a thousand physical qubits per encoded qubit when using traditional surface code frameworks [23], [48].

To reduce this demanding overhead, alternative approaches based on quantum LDPC codes are under active investigation. These codes offer a potential order-of-magnitude reduction in physical qubit overheads [49]. Their realization, however, presents new challenges, such as implementing complex check operators that often require long-range or highly connected qubit layouts [48], [49].

At the moment, most large-scale resource estimates for fault-tolerant algorithms still rely on simplified noise models and optimistic assumptions about control stability. Experimental demonstrations of LDPC codes and other advanced schemes are only beginning to appear and typically involve a small number of qubits. A clear picture of how these codes behave under realistic hardware imperfections is still emerging, which means that projected overhead reductions should be viewed as indicative rather than definitive.

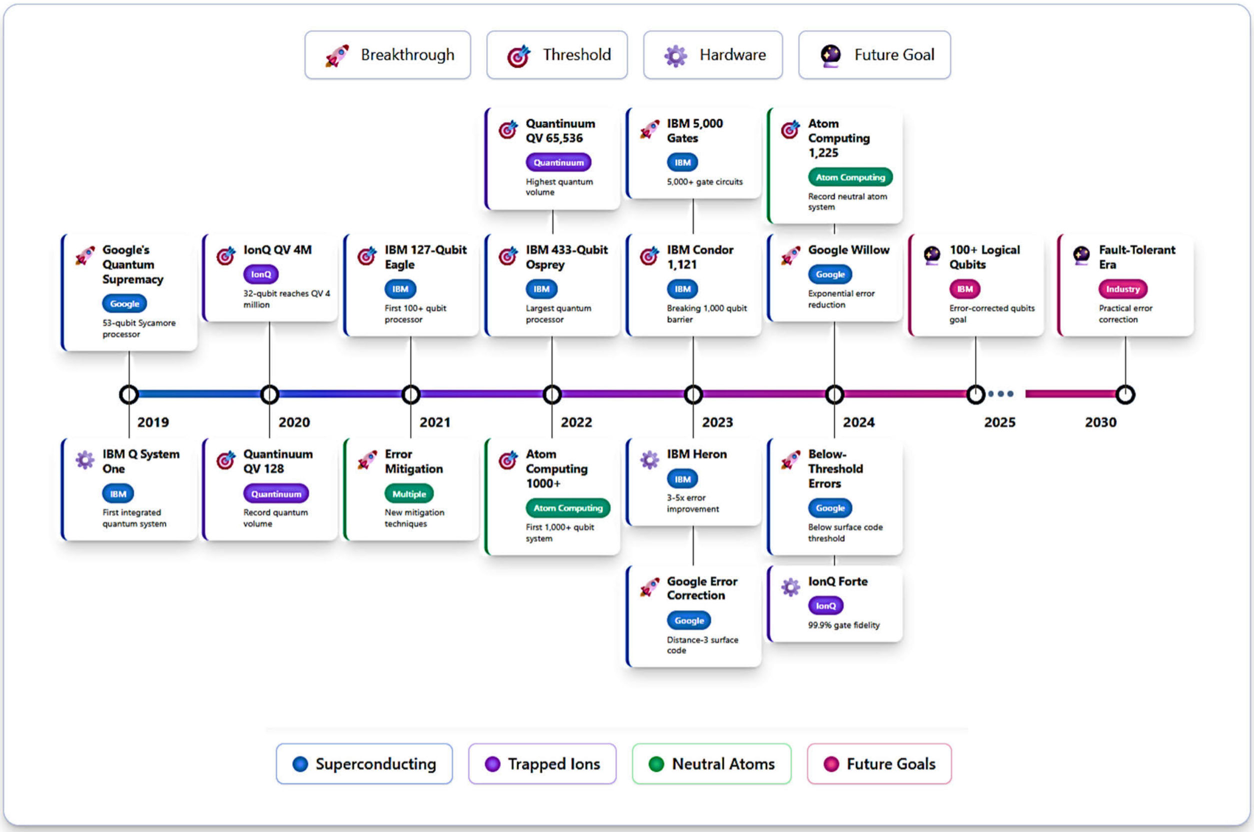

FIGURE 4. Timeline of quantum error-correction milestones(2020–2030), generated using claude 3.5(Anthropic, 2025).

Importantly, fault tolerance is not a binary achievement but a layered continuum. IBM's public roadmap outlines milestones like prototype logical qubits around 2025 and fully error-corrected 1,000-qubit systems by 2030 [7], [48]. Even modest logical arrays, 10 to 50 high-fidelity logical qubits, could unlock valuable applications in fields like quantum chemistry or optimization through mid-depth algorithms such as quantum phase estimation [25].

Consequently, the notion of a "timeline" to fault tolerance is best understood as a sequence of technical thresholds rather than a fixed date. Moving from single logical qubits to small logical registers, and then to algorithmically useful logical systems, is likely to involve multiple iterations of code design, control hardware, and fabrication processes.

Progress toward fault-tolerant quantum computing is expected to depend on the gradual refinement of multiple technical elements. Enhancements in qubit coherence, more reliable control schemes, and scalable error-handling approaches are among the areas receiving attention. Rather than a single breakthrough, it is the accumulation of smaller advancements that may enable quantum systems to support longer and more intricate computational tasks over time.

D. OUTLOOK AND POTENTIAL APPLICATIONS

Although fully fault-tolerant quantum systems remain a distant goal, it is increasingly relevant to consider the capabilities of existing and upcoming devices. Exploring their practical potential highlights opportunities for meaningful applications, even before large-scale error correction becomes viable.

As quantum hardware evolves, identifying meaningful applications for early- and intermediate-scale devices becomes increasingly important. Although the long-term vision involves fully fault-tolerant quantum computers capable of implementing algorithms like Shor's for cryptographic tasks or Grover's for large unstructured searches, such applications are likely to remain out of reach until machines can reliably support thousands of logical qubits [3], [23]. Nonetheless, promising use cases are emerging that do not require full error correction.

Near-term devices, especially those within the noisy intermediate-scale quantum (NISQ) regime, have already begun demonstrating potential in fields such as quantum chemistry, optimization, and machine learning. For instance, variational quantum eigensolver (VQE) techniques have enabled simulations of simple molecules, including hydrogen and lithium hydride, on quantum processors with fewer than 20 qubits [1], [25]. Initial work in quantum chemistry using variational techniques has provided a foundation for exploring applications in computational molecular science. These efforts suggest that quantum processors, even at limited scale, may assist in tasks such as identifying promising drug candidates, analyzing catalytic activity, or characterizing functional materials, especially when paired with conventional simulation tools.

However, most of these demonstrations are still proof-of-concept studies on small problem instances that can also be solved classically. In many cases, classical heuristics or tensor-network methods match or exceed the performance of current quantum approaches at comparable scales. As a result, claims of near-term "quantum advantage" in these domains must be interpreted carefully and are often contingent on very specific problem structures and performance metrics.

In the domain of combinatorial optimization, hybrid quantum-classical strategies such as QAOA have been applied to structured problem instances. These include test cases from logistics, scheduling, and finance, where improvements in solution quality or computational efficiency may emerge as devices improve [29]. Machine learning is also an active frontier, with quantum kernels and hybrid classifiers showing early indications of improved performance on certain structured datasets [25].

Quantum simulation of many-body systems is regarded as one of the most natural and impactful near-term applications. Analog simulators built from ultracold atoms, trapped ions, and Rydberg arrays have already produced insights into phenomena such as quantum thermalization, localization, and phase transitions; regimes that are intractable on classical supercomputers [18], [26]. As these platforms become increasingly tunable and digitized, they may bridge the gap toward programmable quantum simulators capable of modeling exotic materials, superconductivity mechanisms, or high-energy physics scenarios [36].

Quantum-enhanced sensing and metrology represent another domain of growing relevance. Certain quantum processors can be configured to act as ultra-sensitive detectors by leveraging entanglement to improve measurement precision. Trapped ions and nitrogen-vacancy (NV) centers, originally developed for quantum computing, are being repurposed in spectroscopy, navigation, and biomagnetic field detection [24]. Future systems may combine computational and sensing capabilities in hybrid configurations [17].

In tandem with hardware, software techniques such as error mitigation and hybrid computing are being refined to extract performance from imperfect systems. Algorithms that adapt to hardware-specific noise characteristics, such as noise-aware compilation or algorithmic error suppression, can prolong coherence or reduce error accumulation without requiring full-scale error correction [16]. Hybrid classical-quantum strategies remain central: algorithms like VQE and QAOA rely on classical optimizers to guide quantum circuits, creating flexible workflows well suited to the capabilities of current hardware [25].

For combinatorial optimization and machine learning, benchmarking is further complicated by the choice of objective function and classical baseline. Improvements reported on narrow benchmarks may not always translate into practical gains in real-world workflows. Current studies are therefore increasingly focused on identifying problem families where quantum circuits can exploit structure that is hard for classical methods to capture efficiently.

Looking across the ecosystem, the field is undergoing convergence and diversification simultaneously. While some companies aim to scale superconducting or ion-based processors beyond 1,000 physical qubits [7], [17], others explore novel modalities such as fluxonium qubits, photonic architectures, or hybrid approaches that combine different qubit types within a single system [33], [49]. Innovations in classical domains, such as machine learning-guided pulse calibration or lithographic techniques from semiconductor manufacturing, are now being applied to optimize quantum control and fabrication processes [38].

Ultimately, the diversity of hardware paths may itself become an asset. It remains possible that different architectures will find complementary roles: fast superconducting qubits for general-purpose computation, trapped ions for precision tasks, and photonics for scalable communication and networking [6], [13], [24]. Although it remains difficult to predict exactly when large-scale fault-tolerant machines will become viable, incremental gains in device stability, control precision, and software orchestration are steadily extending the scope of addressable quantum problems.

In closing, quantum computing is gradually evolving from a primarily experimental pursuit into a field with clearly defined technical benchmarks and application targets. As research disciplines continue to converge, spanning physics, engineering, computer science, and applied mathematics, it becomes increasingly plausible that quantum systems will begin solving specific, classically intractable problems within the next decade [50].

At the same time, there remains substantial uncertainty about when such problem instances will emerge in practice and how broad their impact will be. Hardware progress, algorithmic innovation, and classical computing advances all move in parallel, and history shows that classical methods often improve in response to proposed quantum speedups. For that reason, most current expectations about near-term utility should be viewed as scenario-based projections rather than firm forecasts.

V. CONCLUSION

The development of quantum processors has become a deeply interdisciplinary endeavor, shaped by advances across physics, engineering, materials science, and software systems. No single qubit technology has emerged as universally dominant; instead, superconducting circuits, trapped ions, photonic systems, neutral atoms, silicon spin qubits, and exploratory topological approaches each exhibit distinct strengths and limitations. This diversity reflects a field that is progressing through parallel technological paths rather than converging prematurely on a single architecture.

Current evidence suggests that scalable quantum computing is unlikely to be realized through isolated hardware improvements alone. Instead, hybrid and modular system designs, combining complementary qubit platforms with classical co-processing are emerging as a practical strategy for near- and mid-term progress. In this context, the concept of quantum advantage is evolving from isolated benchmark demonstrations toward domain-specific quantum utility, where carefully matched algorithms and hardware can deliver value despite the constraints of noisy intermediate-scale devices. Early opportunities are therefore most plausible in structured problem classes such as quantum simulation, materials modeling, and specialized optimization tasks.

Equally important is the coordinated orchestration of hybrid quantum–classical systems across the full computing stack, where advances in hardware, control electronics, calibration, compilation, and algorithms increasingly determine real-world quantum performance. Architectural co-design and hardware-aware software are therefore becoming essential for extracting usable performance, often outweighing incremental gains in individual device metrics. As quantum cloud platforms mature, interoperability and flexible software toolchains are also expected to play a growing role in enabling users to access and combine heterogeneous hardware resources effectively.

Looking forward, the most promising path to meaningful quantum impact is not linear scaling toward fault tolerance alone, but strategic convergence: modular hardware architectures, noise-tolerant algorithmic techniques, and application domains that reward structured quantum heuristics. In this sense, the present diversity of quantum platforms is a strength rather than a liability, providing a broad experimental foundation from which robust, scalable systems can emerge. While fully fault-tolerant quantum computers remain a longer-term goal, recent experimental milestones and application-driven demonstrations increasingly suggest that the field is approaching a transition from milestone advantage demonstrations toward reproducible and application-driven quantum utility [17], [28]. In this context, the central question is shifting from whether quantum computing will have an impact to where and how such impact is most likely to materialize first, typically through incremental, system-level advances rather than a single disruptive breakthrough. As benchmarking methodologies mature and logical qubit demonstrations expand, clearer comparisons between NISQ performance and fault-tolerant scalability are expected to guide the evolution of quantum processor architectures.

REFERENCES

[1] M. Benedetti, M. Fiorentini, and M. Lubasch, "Hardware-efficient variational quantum algorithms for time evolution," *Phys. Rev. Res.*, vol. 3, no. 3, Jul. 2021.

[2] R. P. Feynman, "Simulating physics with computers," *Int. J. Theor. Phys.*, vol. 21, nos. 6–7, pp. 467–488, Jun. 1982.

[3] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, U.K.: Cambridge Univ. Press, 2010.

[4] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, Aug. 2018.

[5] F. Arute et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.

[6] M. H. Devoret and R. J. Schoelkopf, "Superconducting circuits for quantum information: An outlook," *Science*, vol. 339, no. 6124, pp. 1169–1174, Mar. 2013.

[7] (2024). *Development & Innovation Roadmap*. [Online]. Available: https://www.ibm.com/quantum/assets/IBM_Quantum_Developmen_&_Innovation_Roadmap_Explainer_2024-Update.pdf

[8] M. AbuGhanem, "Superconducting quantum computers: Who is leading the future?" *EPJ Quantum Technol.*, vol. 12, no. 1, Dec. 2025.

[9] R. Blatt and D. Wineland, "Entangled states of trapped atomic ions," *Nature*, vol. 453, no. 7198, pp. 1008–1015, 2008.

[10] K. Wright et al., "Benchmarking an 11-qubit quantum computer," *Nature Commun.*, vol. 10, no. 1, p. 5464, 5464.

[11] J. M. Pino et al., "Demonstration of the trapped-ion quantum CCD computer architecture," *Nature*, vol. 592, no. 7853, pp. 209–213, 2021.

[12] (2025). *GovInfo*. Accessed: Jun. 1, 2025. [Online]. Available: https://www.govinfo.gov/app/details/PLAW-115publ368

[13] J. L. O'brien, A. Furusawa, and J. Vuckovic, "Photonic quantum technologies," *Nature Photon.*, vol. 3, no. 12, pp. 687–695, 2009.

[14] M. Thompson, "Photonic quantum computing," in *Proc. 28th Microoptics Conf. (MOC)*, Miyazaki, Japan, Sep. 2023, pp. 1–2, doi: 10.23919/moc58607.2023.10302854.

[15] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, "Validating quantum computers using randomized model circuits," *Phys. Rev. A, Gen. Phys.*, vol. 100, no. 3, Sep. 2019.

[16] Z. Cai et al., "Quantum error mitigation," *Rev. Mod. Phys.*, vol. 95, no. 4, 2023, Art. no. 045005.

[17] Y. Kim et al., "Evidence for the utility of quantum computing before fault tolerance," *Nature*, vol. 618, no. 7965, pp. 500–505, 2023.

[18] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, "Experimental issues in coherent quantum-state manipulation of trapped atomic ions," *J. Res. Nat. Inst. Standards Technol.*, vol. 103, no. 3, p. 259, May 1998.

[19] S. A. Moses et al., "A race-track trapped-ion quantum processor," *Phys. Rev. X*, vol. 13, no. 4, 2023, Art. no. 041052.

[20] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, "Superconducting qubits: Current state of play," *Annu. Rev. Condens. Matter Phys.*, vol. 11, no. 1, pp. 369–395, Mar. 2020.

[21] Y. Shapira, J. Markov, N. Akerman, A. Stern, and R. Ozeri, "Programmable quantum simulations on a trapped-ion quantum computer with a global drive," *Phys. Rev. Lett.*, vol. 134, no. 1, Jan. 2025.

[22] A. D. King et al., "Observation of topological phenomena in a programmable lattice of 1,800 qubits," *Nature*, vol. 560, no. 7719, pp. 456–460, Aug. 2018.

[23] V. E. Elfving, B. W. Broer, M. Webber, J. Gavartin, M. D. Halls, K. P. Lorton, and A. Bochevarov, "How will quantum computers provide an industrially relevant computational advantage in quantum chemistry?" 2020, *arXiv:2009.12472*.

[24] D. D. Awschalom, R. Hanson, J. Wrachtrup, and B. B. Zhou, "Quantum technologies with optically interfaced solid-state spins," *Nature Photon.*, vol. 12, no. 9, pp. 516–527, 2018.

[25] K. Bharti et al., "Noisy intermediate-scale quantum algorithms," *Rev. Mod. Phys.*, vol. 94, no. 1, 2022, Art. no. 015004.

[26] X. Wu, X. Liang, Y. Tian, F. Yang, C. Chen, Y. C. Liu, M. K. Tey, and L. You, "A concise review of Rydberg atom-based quantum computation and quantum simulation," *Chin. Phys. B*, vol. 30, no. 2, 2021, Art. no. 020305.

[27] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, "High-threshold and low-overhead fault-tolerant quantum memory," *Nature*, vol. 627, no. 8005, pp. 778–782, Mar. 2024.

[28] R. Acharya et al., "Suppressing quantum errors by scaling a surface code logical qubit," *Nature*, vol. 614, pp. 676–681, Jan. 2023.

[29] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014, *arXiv:1411.4028*.

[30] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc.28th Annu. ACM Symp. Theory Comput. - STOC*, 1996, pp. 212–219.

[31] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, Sep. 2017.

[32] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," *Nature Commun.*, vol. 5, no. 1, Jul. 2014.

[33] N. C. Harris, G. R. Steinbrecher, M. Prabhu, Y. Lahini, J. Mower, D. Bunandar, C. Chen, F. N. C. Wong, T. Baehr-Jones, M. Hochberg, S. Lloyd, and D. Englund, "Quantum transport simulations in a programmable nanophotonic processor," *Nature Photon.*, vol. 11, no. 7, pp. 447–452, Jul. 2017.

[34] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, A. E. Lita, T. Gerrits, S. W. Nam, V. D. Vaidya, M. Menotti, I. Dhand, Z. Vernon, N. Quesada, and J. Lavoie, "Quantum computational advantage with a programmable photonic processor," *Nature*, vol. 606, no. 7912, pp. 75–81, Jun. 2022.

[35] M. A. Serrano et al., "Quantum software components and platforms: Overview and quality assessment," *ACM Comput. Surv.*, vol. 55, no. 8, pp. 1–31, 2022.

[36] M. Saffman, T. G. Walker, and K. Mølmer, "Quantum information with Rydberg atoms," *Rev. Modern Phys.*, vol. 82, no. 3, pp. 2313–2363, Aug. 2010.

[37] A. Omran et al., "Generation and manipulation of Schrödinger cat states in Rydberg atom arrays," *Science*, vol. 365, no. 6453, pp. 570–574, 2019.

[38] W. D. Oliver and P. B. Welander, "Materials in superconducting quantum bits," *MRS Bull.*, vol. 38, no. 10, pp. 816–825, Oct. 2013.

[39] N. Sachdeva et al., "Quantum optimization using a 127-qubit gate-model IBM quantum computer can outperform quantum annealers for nontrivial binary optimization problems," 2024, *arxiv: 2406.01743v1*.

[40] F. A. Zwanenburg et al., "Silicon quantum electronics," *Rev. Mod. Phys.*, vol. 85, no. 3, pp. 961–1019, 2013.

[41] T. Meunier, N. Daval, F. Perruchot, and M. Vinet, "Silicon spin qubits: A viable path towards industrial manufacturing of large-scale quantum processors," *Eur. Phys. J. A*, vol. 61, no. 3, pp. 1–11, Mar. 2025.

[42] F. Borjans et al., "Resonant microwave-mediated interactions between distant electron spins," *Nature*, vol. 577, no. 7789, pp. 195–198, 2020.

[43] M. Veldhorst, C. H. Yang, J. C. C. Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak, "A two-qubit logic gate in silicon," *Nature*, vol. 526, no. 7573, pp. 410–414, Oct. 2015.

[44] A. Y. Kitaev, "Fault-tolerant quantum computation by anyons," *Ann. Phys.*, vol. 303, no. 1, pp. 2–30, Jan. 2003.

[45] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. D. Sarma, "Non-abelian anyons and topological quantum computation," *Rev. Mod. Phys.*, vol. 80, no. 3, pp. 1083–1159, 2008.

[46] M. Aghaee et al., "InAs-Al hybrid devices passing the topological gap protocol," *Phys. Rev. B, Condens. Matter*, vol. 107, no. 24, 2023, Art. no. 245423.

[47] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Phys. Rev. A, Gen. Phys.*, vol. 86, no. 3, Sep. 2012.

[48] P. Panteleev and G. Kalachev, "Quantum LDPC codes with almost linear minimum distance," *IEEE Trans. Inf. Theory*, vol. 68, no. 1, pp. 213–229, Jan. 2022.

[49] C. Ryan-Anderson et al., "Realization of real-time fault-tolerant quantum error correction," *Phys. Rev. X*, vol. 11, no. 4, 2021, Art. no. 041058.

[50] S. Lloyd, "Ultimate physical limits to computation," *Nature*, vol. 406, no. 6799, pp. 1047–1054, Aug. 2000.

MARKO VUKŠIĆ received the M.S. degree in electronic and information technologies in maritime studies from the University of Rijeka, Rijeka, Croatia, where he is currently pursuing the Ph.D. degree in maritime studies with the Department of Electrical Engineering, Automation and Computing, Faculty of Maritime Studies. He has participated in research projects related to cybersecurity, resilience, and sustainability in the maritime industry and has published in international journals and conferences. His research interests include maritime cybersecurity, cyber risk assessment, and resilient shipboard information systems.

JASMIN ĆELIĆ (Senior Member, IEEE) received the Ph.D. degree in maritime studies from the Department of Electrical Engineering, Automation and Computing, Faculty of Maritime Studies, University of Rijeka, Rijeka, Croatia. He is currently an Associate Professor with the Department of Electrical Engineering, Automation and Computing, Faculty of Maritime Studies, University of Rijeka. He has led and contributed to numerous national and international research projects in maritime cybersecurity, intelligent transport systems, and cyber-resilient marine technologies. His research has been published in indexed journals and international conferences. His current research interests include maritime cybersecurity, intelligent transportation systems, and smart mobility infrastructures.

ALEKSANDAR CUCULIĆ received the Ph.D. degree in maritime studies from the Department of Electrical Engineering, Automation and Computing, Faculty of Maritime Studies, University of Rijeka, Rijeka, Croatia. He is currently an Associate Professor with the Department of Electrical Engineering, Automation and Computing, Faculty of Maritime Studies, University of Rijeka. He has led and contributed to national and international research projects in maritime energy systems, sustainable vessel technologies, and intelligent port infrastructures. His research has been published in indexed journals and presented at international conferences. His current research interests include marine electrical power systems, hybrid and electric propulsion solutions, maritime energy resilience, and smart port energy technologies.

• • •