State of the Art and Roadmap to a Million-Qubit Quantum Computer (DRAFT)

Paul F. Roysdon, Ph.D.

Published: March 3, 2025, Updated: February 22, 2026

Abstract

Quantum computing promises to tackle problems intractable for classical machines by exploiting quantum phenomena such as superposition and entanglement. Planning estimates for utility-scale fault-tolerant quantum computing increasingly span **10^5–10^6 physical qubits**, depending on the error-correcting code family, decoder latency, and modular architecture. Depending on code and architectural assumptions, recent resource analyses suggest RSA-2048 could require on the order of 10^5–10^6 physical qubits for week-scale runtimes, motivating careful algorithm–QEC co-design. This paper provides a comprehensive roadmap and survey of the state-of-the-art toward building a million-qubit quantum computer. We integrate advances in hardware, software, and theory – detailing the physics of leading qubit implementations and the latest error-correction techniques – and present a systems engineering perspective on scaling quantum processors to the million-qubit level. We outline a stepwise development plan toward this goal and review major public and private initiatives supporting the effort. The aim is to chart the path from today's intermediate-scale prototypes to tomorrow's fault-tolerant quantum supercomputers, highlighting both the immense challenges and the profound opportunities of the quantum computing revolution.

I. Current State of the Art in Quantum Computing

Quantum computing has rapidly advanced in recent years, with progress on multiple hardware platforms, software frameworks, and theoretical foundations. This section surveys the state of the art in leading qubit technologies, quantum software and benchmarking tools, and key theoretical advances such as quantum error correction and algorithm optimization.

A. Hardware Platforms for Qubits

Superconducting Qubits (Transmons): Superconducting qubits are among the most developed platforms, used by IBM, Google, and others. They consist of superconducting circuits incorporating Josephson junctions, cooled to millikelvin temperatures to exhibit quantum effects. Superconducting devices have demonstrated chip-scale qubit counts into the hundreds and beyond (e.g., IBM's 433-qubit "Osprey" and 1121-qubit "Condor" announcements), while recent roadmaps emphasize improving fidelity and modular scaling via system-level integration (e.g., IBM Quantum System Two built around multiple "Heron" processors) [3], [26]. Coherence times and two-qubit fidelities vary by design and operating point: best-in-class demonstrations can approach millisecond-scale T_1 and $\sim$99.9% two-qubit fidelity in carefully tuned settings, whereas scalable, densely integrated processors often operate in the tens-to-hundreds of microseconds regime with device-dependent two-qubit error rates. An advantage of superconducting qubits is their fast operation speed – gate times are on the order of tens of nanoseconds, making them about 1000$\times$ faster than gates in ion or atom systems [3]. However, superconducting qubits have relatively limited connectivity (typically coupling only to nearest neighbors on a chip) and require complex cryogenic infrastructure. Major industry milestones include IBM's 127-qubit and 433-qubit processors and the 1121-qubit "Condor" chip delivered in 2023 [3], as well as Google's 53-qubit Sycamore processor which achieved a quantum supremacy demonstration in 2019 [12].

Trapped Ion Qubits: Trapped ion quantum computers use individual ions (charged atoms) confined in electromagnetic traps under ultra-high vacuum. The ions (e.g., Yb^+, Ca^+) are laser-cooled to their motional ground state and encode qubits in internal electronic states (such as hyperfine or metastable states) [8]. Laser beams or microwave fields perform quantum gate operations by coupling to the ions' electronic and motional states [8]. A key advantage of trapped ions is that all qubits of the same species are identical and can interact via collective motion, enabling inherently high-fidelity gates and all-to-all connectivity within a trap [8]. Single- and two-qubit gate fidelities in trapped-ion systems routinely exceed 99% [8], and coherence times are very long (ions can maintain quantum states for tens of seconds or more). For example, quantum gates, state preparation, and measurement in ion traps have achieved fidelities above 99% [8]. Trapped ions are thus highly reliable qubits. The trade-off is slower gate speeds (two-qubit gates often take tens of microseconds or more [8]) and challenges in scaling to large numbers of ions in a single trap. Current ion processors contain on the order of 10-50 qubits in a linear chain. Companies like IonQ and Quantinuum (Honeywell) have demonstrated systems in this range with record high fidelities. Research toward scaling includes the quantum charge-coupled device (QCCD) architecture, where ions are shuttled between trap zones [8], and modular ion trap networks connected by photonic links [8]. These approaches aim to overcome the difficulty of controlling very long ion chains by networking smaller ion modules. Despite modest qubit numbers, trapped-ion systems have shown quantum volume and algorithmic performance on par with or better than other platforms due to their precision. They have been used to run algorithms in chemistry, optimization (e.g., Quantum Approximate Optimization Algorithm [QAOA]), and machine learning, demonstrating the platform's versatility [8].

Neutral Atom Qubits: Neutral atom quantum computers use neutral (uncharged) atoms trapped in arrays by highly focused laser beams called optical tweezers. Atoms such as rubidium or cesium are first cooled (e.g., via magneto-optical traps and Doppler cooling) to micro-Kelvin temperatures to nearly freeze out thermal motion [9]. Individual atoms are then captured and arranged into a regular 2D (or 3D) array using optical tweezers, with feedback to rearrange atoms into filled lattice sites [9]. Qubits are encoded in atomic states, typically the ground state $|0\rangle$ and a Rydberg excited state or another long-lived state $|1\rangle$. Two-qubit gates are implemented via the Rydberg blockade mechanism: exciting one atom to a high-lying Rydberg state induces a strong dipole-dipole interaction that shifts the energy of a neighboring atom, preventing its excitation [9]. This controlled blockade serves as an entangling gate between atoms. Neutral atom systems offer flexible, reconfigurable qubit connectivity – atoms can be arranged in arbitrary geometries, and long-range Rydberg interactions allow entangling non-nearest neighbors [9]. They also feature long coherence times (qubit states can remain coherent for seconds, far longer than typical gate durations) [9]. Recent neutral-atom hardware has achieved arrays of $\sim$100 or more atoms, with plans to scale to 1000 qubits in coming years [9]. Gate fidelities have steadily improved, now around $\sim$99% for two-qubit gates and $>$99.5% for single-qubit operations [9]- approaching thresholds for quantum error correction. While two-qubit gates are still somewhat slower (microseconds range) and less error-free than in the best ion/superconducting systems, the rapid progress is closing this gap. Notably, QuEra Computing demonstrated a 256-atom quantum simulator (for analog Hamiltonian simulation), and startups like Atom Computing and Pasqal are developing digital gate-based neutral atom processors. The ability to scale to 2D/3D atom arrays and the prospect of photonic integrated control make neutral atoms a promising platform for large-scale quantum computers [9].

Photonic Qubits: Photonic quantum computing uses individual photons as carriers of quantum information. A photonic qubit can be encoded in various photon degrees of freedom – common encodings include polarization (horizontal/vertical), optical path modes (dual-rail), time-bin (early vs. late arrival), or frequency [10]. Photons have the great advantage of very low decoherence: they do not easily interact with the environment, so quantum states of light can in principle travel long distances or persist for long times with minimal noise [10]. Moreover, photonic qubits naturally operate at room temperature and can leverage existing telecom fiber and integrated optics infrastructure [10]. The biggest challenge is that photons do not naturally interact with each other,

making two-qubit gates difficult. Linear optical approaches (using beam splitters and phase shifters) combined with measurements can produce probabilistic entangling gates, but deterministic entanglement usually requires special resource states or nonlinear interactions. One leading approach is **measurement-based quantum computing** using cluster states of many entangled photons: a large entangled photonic graph state is generated, and then performing adaptive measurements on subsets of photons effectively implements quantum logic on the remaining photons [9]. Companies like PsiQuantum and Xanadu are pursuing photonic quantum computers by creating on-chip light sources and detectors to generate massive entangled cluster states for fault-tolerant computation. To this end, significant work is ongoing in developing deterministic single-photon sources (e.g., quantum dot emitters, parametric down-conversion sources) and high-efficiency single-photon detectors [10]. Progress is being made: researchers have demonstrated photonic processors capable of specific tasks like boson sampling with $>$100 photons (an approach to show quantum advantage). PsiQuantum, for example, is partnering with GlobalFoundries to fabricate silicon photonic chips and superconducting photon detectors at scale [16]. While no universal photonic gate-model quantum computer exists yet, the photonic platform is considered scalable in the long run because it can potentially integrate on semiconductor chips and because photons can be easily routed and duplicated. Photonic qubits are also crucial for quantum networking and communication. In summary, photonic quantum computing is currently at a prototype stage, but its long coherence and room-temperature operation make it an important part of the quantum technology landscape [9], [10].

Spin Qubits in Semiconductors: Spin qubits leverage the quantum spin of electrons or nuclei in solid-state devices (most commonly, electrons in silicon-based quantum dots or impurity atoms). These are often referred to as **silicon spin qubits**, and they have the attractive feature of compatibility with standard semiconductor fabrication – essentially, tiny quantum dot "transistors" can confine single electrons whose spin-up vs. spin-down serves as $|0\rangle/|1\rangle$ [3]. Because they are built in silicon, spin qubits are extremely small (on the order of 50-100 nm) and could be densely integrated by leveraging Complementary Metal-Oxide-Semiconductor (CMOS) manufacturing techniques [3]. Recent advances show that silicon spin qubits have excellent coherence properties: through isotopic purification of silicon (removing nuclear spin noise) and careful materials engineering, coherence times up to 0.5 seconds have been achieved [11] – exceptionally long for a solid-state qubit. High-fidelity control has also been demonstrated, with single-qubit gate fidelities $>$99.95% and two-qubit gate fidelities surpassing the fault-tolerance threshold (e.g., $\sim$99.9%) [11]. Moreover, these fidelities have even been reached at temperatures of 1-4 K ("hot" by quantum standards), raising the possibility of operating spin qubits with simpler cryogenics [11]. Basic one- and two-qubit logic is implemented via electron spin resonance (microwave pulses to flip spins) and exchange coupling between adjacent spins in a two-electron quantum dot (enabling CNOT, SWAP, etc.) [11]. There are two main implementations: gate-defined quantum dots (electrostatically confining electrons in a 2D electron gas) and donor spins (using dopant atoms like phosphorus in silicon) [11]. Both approaches have shown success; for instance, lab prototypes have entangled up to 6-10 spins in linear arrays, and efforts are underway to demonstrate larger 2D spin qubit arrays. A unique advantage of spin qubits is the potential to leverage semiconductor industry techniques for scaling – billions of classical transistors are routinely made on a chip, so in principle millions of quantum dots might be fabricated. Intel and research partners have already developed 300 mm wafer processes for spin qubits with high uniformity [3]. Spin qubits are still in the research stage relative to superconducting or ion qubits, with smaller numbers demonstrated so far, but their long coherence and scaling promise make them a leading candidate for future large-scale quantum processors [3]. Indeed, only spin qubits and superconducting qubits presently seem amenable to monolithic integration of millions of units, and spin qubits' small size and fabrication compatibility give them a scaling edge [3]. Recent milestones include strong spin-photon coupling (using superconducting resonators to link distant spins) [11]and the development of cryogenic CMOS control chips to interface with spin qubit arrays. In summary, silicon spin qubits are emerging as a highly promising platform combining the long coherence of isolated atomic systems with the manufacturability of solid-state

technology [11].

To summarize, Table I compares key metrics of major qubit platforms, highlighting their performance and scalability:

TABLE I

COMPARISON OF QUBIT HARDWARE PLATFORMS (APPROXIMATE RANGES; BEST-IN-CLASS DEMONSTRATIONS AND "TYPICAL" FLEET PERFORMANCE CAN DIFFER SUBSTANTIALLY).

Qubit Platform	Coherence Time	Two-Qubit Gate Fidelity	Gate Speed (approx.)	Connectivity & Scale
Superconducting (transmon)	$\sim$50–300 μs typical (best approaching $\sim$ms)	device-dependent; best reports approach $\sim$99.9%	$\sim$20–100 ns (GHz-scale gates)	Planar circuit, nearest-neighbor coupling; 100s–1000+ qubits demonstrated; modular scaling under active development.
Trapped Ion	$\sim$10–100 s (very long)	$>$99% (best reports $\sim$99.9%+)	$\sim$10–100 μs (laser-mediated)	All-to-all coupling in small chains; tens of ion qubits (50+ reported) with very high per-qubit quality.
Neutral Atom (Rydberg)	$\sim$1–5 s (long-lived states)	$\sim$99% (device- and protocol-dependent)	$\sim$1–5 μs (Rydberg blockade)	Flexible 2D/3D connectivity via Rydberg links; 100s to 1000+ atom arrays demonstrated (often analog), with digital FT roadmaps in progress.
Photonic (optical)	Essentially unlimited in flight (no decoherence)	N/A (entangling operations via measurements)	Fast (at light speed; GHz detectors)	Linear optics gives limited direct connectivity; large entangled cluster states ($>$100 photons) produced in experiments; FT architectures in development.
Semiconductor Spin (silicon)	$\sim$0.1–0.5 s (with isotopic purification)	$\sim$99.9% (selected demonstrations)	$\sim$10–100 ns (spin resonance)	Local coupling (exchange or via resonators); small arrays (single- to tens-of-qubits) so far, but CMOS fabrication offers a path to very large scale.

B. Software Frameworks and Benchmarking Tools

Accompanying the hardware advances is a rich ecosystem of quantum software and benchmarks to program and evaluate quantum computers. Several open-source frameworks have become standard for writing quantum algorithms and controlling quantum hardware:

- **Qiskit (IBM)**: A comprehensive software stack in Python for writing quantum programs at the circuit level, running them on IBM quantum processors or simulators, and performing higher-level algorithm

development. It includes modules for algorithm libraries, pulse-level control, and tools for error mitigation and benchmarking.

- **Cirq (Google)**: A Python framework for creating and optimizing quantum circuits, especially suited to NISQ devices. Cirq gives users fine-grained control over gate operations and has been used in Google's quantum supremacy and error-correction experiments.
- **PyQuil (Rigetti) and Forest SDK**: Tools for programming Rigetti's superconducting qubit systems, featuring a Quil instruction language and a hybrid quantum/classical programming model.
- **Q#**: Microsoft's quantum programming language, integrated with Visual Studio and Azure Quantum, focused on quantum algorithm development and resource estimation. Q# is paired with a simulator and tools to estimate the number of qubits/gates needed for quantum algorithms, which is crucial for long-term algorithm design.
- **PennyLane, Cirq, and Others for Hybrid Algorithms**: Frameworks like PennyLane and TensorFlow Quantum allow integration of quantum circuits as layers in classical machine learning models, enabling variational quantum algorithms and quantum neural networks on various hardware backends.

To assess and compare quantum hardware performance, several **benchmarking metrics** have emerged:

- **Quantum Volume (QV)**: A holistic benchmark introduced by IBM that measures the largest random circuit of equal width and depth that a quantum device can successfully implement. A higher QV indicates better overall qubit count, connectivity, and fidelity. IBM has steadily increased QV for its devices (doubling nearly yearly); for instance, IBM reported achieving Quantum Volume 256 (2^8) on a 27-qubit device in 2021 [19]. QV accounts for crosstalk and gate errors, providing a single-number metric of "quantum computational power."
- **CLOPS (Circuit Layer Operations Per Second)**: IBM also introduced CLOPS to measure the execution speed of quantum programs (how many layers of gates per second a system can perform, including classical control overhead). This addresses not just fidelity but also how fast a quantum processor can be driven, which is important for algorithms requiring many circuit iterations [19].
- **Randomized Benchmarking and Error Rates**: At a lower level, quantum gates are characterized by randomized benchmarking, which yields error rates per gate (e.g., 1-qubit gate error $\sim 10^{-4}$, 2-qubit $\sim 10^{-3}$ for top devices [3]). These error rates feed into error correction thresholds.
- **Algorithmic Qubits and Quantum Advantage Benchmarks**: Companies also propose problem-centric benchmarks. For example, Quantinuum uses an "**algorithmic qubit**" metric indicating the largest problem size (like quantum circuit complexity) the system can handle with high fidelity. Google demonstrated a beyond-classical **random circuit sampling** on 53 qubits [12], and USTC in China demonstrated photonic boson sampling with 113 photons – these are specialized benchmark tasks to claim quantum advantage.
- **Application-specific benchmarks**: e.g., simulating a specific molecule's energy or solving an optimization problem to a certain quality. These end-to-end benchmarks are increasingly important to measure progress toward useful quantum advantage in areas like chemistry or machine learning.

The quantum software stack also includes **error mitigation and compilation** tools that squeeze the best performance out of hardware. Techniques such as transpilation (reordering and optimizing gates to fit hardware constraints), dynamical decoupling (inserting pulses to combat decoherence), zero-noise extrapolation and probabilistic error cancellation (to reduce error effects without full error correction) are implemented at the software level to maximize algorithm success on today's noisy devices.

In summary, the community has developed robust software frameworks to program quantum computers and a variety of benchmarks to measure progress. These tools are crucial for evaluating improvements as devices scale up. For instance, IBM's **1000+ qubit Condor processor** will be assessed not just by qubit count but by metrics like quantum volume and error rates to ensure quality scales with quantity [3]. Similarly,

as new algorithms (like variational quantum eigensolvers or Shor's algorithm variants) are tested, software infrastructure is in place to run them and estimate the resources needed for quantum advantage.

C. Theoretical Advances: Error Correction, Algorithms, and Resource Estimation

Alongside hardware and software, there have been significant theoretical advances that shape the path to large-scale quantum computing. Three crucial areas are quantum error correction, algorithmic improvements, and resource estimation for practical quantum advantage.

Quantum Error Correction (QEC): Quantum error correction is essential for achieving fault-tolerant, reliable quantum computation by encoding logical qubits into many physical qubits to detect and correct errors. In theory, the surface code and other QEC codes have an error threshold around $\sim$1% per gate – if physical error rates are below this, arbitrarily long computations are possible with enough redundancy. Recent years have seen both theory and experimental milestones in QEC. Notably, in 2023 Google Quantum AI demonstrated for the first time that increasing the size of a surface code logical qubit (from distance-3 to distance-5, using 49 physical qubits) resulted in lower error rates for the logical qubit, indicating the beginning of the error-correcting regime where quantum error correction actually suppresses error probability [18]. This was a crucial validation of QEC principles in practice. IBM has also pursued quantum error correction aggressively; IBM's 2025 roadmap introduces a milestone of a 1000+ qubit system ("Blue Jay") capable of running 1 billion quantum gate operations using quantum error correction [3]. IBM reported developing new error-correcting code techniques tailored to their hardware that drastically reduce overhead [13]. For example, IBM changed its approach to "first figure out what chips can be built, then design an error correction scheme around those chips," which led to a new algorithm needing far fewer qubits for QEC [13]. Indeed, IBM's Jay Gambetta stated that as of 2025, "**We've answered those science questions [of QEC]. You don't need a miracle now – now it's an engineering challenge.**" [13].

Multiple QEC codes are being explored (surface codes, bosonic codes, error-detecting schemes on ion traps, etc.), but surface codes remain a leading approach for 2D architectures. The **resource overhead** of QEC is still enormous – a logical qubit might require hundreds or thousands of physical qubits. However, improvements are reducing this. For instance, researchers discovered techniques like "**dynamic logical qubits**" or **flag qubits** that reduce overhead, and Google's 2024 work on multi-layer error correction (e.g., a second layer of encoding plus "magic state" distillation optimizations) cut the qubit cost for factoring in one recent estimate [14]. Quantum error correction is transitioning from pure theory to partial practice: small quantum codes (like the [[7,1,3]] Steane code, Shor code, or distance-3 surface code) have been realized on real hardware, and quantum memory lifetimes have been extended via error correction. These successes bolster confidence that a full fault-tolerant quantum computer is feasible with sufficient scale. The **threshold theorem** assures that with error rates below $\sim 10^{-3}$ and reasonable overhead, logical error rates can be exponentially suppressed.

Emerging QEC Direction: qLDPC and Low-Overhead Codes: While the surface code remains a baseline due to its locality and high threshold, recent results have accelerated interest in *quantum low-density parity-check (qLDPC)* codes and related constructions that can reduce logical-qubit overhead (especially for memory) at the cost of more demanding connectivity and decoding. Recent demonstrations of high-threshold, low-overhead fault-tolerant quantum memory in qLDPC families have strengthened the case for architectural co-design around modular interconnects and fast, hardware-friendly decoders [27], [28]. **Algorithmic Optimization (Shor's Algorithm and Beyond)**: On the algorithms front, one famous example is **Shor's algorithm** for factoring large numbers exponentially faster than classical algorithms. Since Shor's 1994 discovery, implementation has been limited to small numbers (due to hardware limits), but ongoing theoretical work focuses on reducing the resources required by Shor's algorithm and related algorithms. A

major 2023 breakthrough by researchers (Chevignard, Fouque, and Schrottenloher) optimized the modular exponentiation step in Shor's algorithm by using an approximate arithmetic technique [14]. Their algorithm dramatically reduced the size of quantum arithmetic circuits, at the cost of more repetitions, which Google's team further optimized down to only $2\times$ overhead [14]. Resource estimates for cryptographically relevant algorithms have tightened considerably and remain sensitive to physical error rates, cycle time, connectivity, and decoder assumptions. A 2025 analysis estimates that factoring an RSA-2048 integer could be achievable in under a week with *fewer than one million* noisy physical qubits using surface-code error correction under explicit, optimistic parameters [14]. More recent architectural proposals that leverage qLDPC codes argue that the physical-qubit requirement could drop into the $\sim 10^5$ range under different assumptions (code family, connectivity, decoder performance, and physical error model) [29]. These results highlight that "a million qubits" is a useful planning scale, but not a universal threshold; algorithm and QEC co-design can shift the required physical resources by orders of magnitude.

Beyond Shor's algorithm, other algorithms are being refined: Grover's search algorithm is being incorporated into hybrid quantum-classical workflows for speedups in unstructured search problems; quantum simulation algorithms for chemistry (like variational quantum eigensolvers) are being improved through better ansatz states and error mitigation; and new algorithms for linear systems, differential equations, and machine learning (quantum ML) are being designed. **Quantum approximate optimization algorithm (QAOA)** is being optimized for specific tasks and hardware. In many cases, the focus is on **noise-tailored algorithms** – algorithms that can give useful results before full error correction by tolerating some noise or leveraging shorter circuits.

In the theoretical computer science domain, researchers are also clarifying the boundary of **quantum advantage** – identifying which problems truly yield exponential vs. polynomial speedups, and which require prohibitively large quantum resources. This guides where to allocate effort. For instance, recent studies on quantum complexity suggest that certain machine learning or optimization tasks have at most polynomial quantum speedups, whereas factoring, quantum simulation of quantum chemistry, and some cryptographic tasks remain the clearest exponential quantum advantages.

Resource Estimation: Hand-in-hand with algorithms and QEC is the field of resource estimation – rigorously determining how many qubits, how many gate operations, and how long it would take for a quantum computer to solve a given problem, under realistic error-corrected conditions. This field has matured significantly. As mentioned, updated resource estimates for breaking RSA encryption have been published by experts at companies like Google [14]. Similarly, estimates for quantum chemistry problems (like accurately simulating a molecule such as FeMoco, a complex bio-catalyst) have been performed – often concluding on the order of a few thousand logical qubits and a few hours of runtime could surpass classical methods for certain chemical problems. These analyses feed into national security and industry planning by indicating when quantum computers become a threat to cryptography or useful for materials science.

One notable concept is the "**logical qubit footprint**" – how many physical qubits and gates are needed per logical qubit. Google's 2025 study highlighted improvements in error-correction that triple the storage density of idle logical qubits by using a second layer of encoding [14]. They also introduced techniques like **magic state cultivation** (an improvement over magic state distillation for generating non-Clifford resources) that reduce the number of ancillary qubits needed [14]. These reduce the overhead for any quantum algorithm, not just factoring. Thanks to such advances, the rough consensus is that a **million-qubit scale** machine with error correction could run some of the most demanding known algorithms (like breaking RSA or simulating complex quantum systems) within days or weeks [14]. This figure – one million physical qubits – has become a *symbolic target* for the quantum computing field, representing the threshold at which a fault-tolerant quantum computer could do impactful, classically intractable tasks [16]. Many roadmaps, as we'll see, explicitly aim

for this scale.

In summary, theory work in QEC is turning error correction from an abstract concept into engineering reality, algorithm research (exemplified by Shor's algorithm optimizations) is cutting down the requirements for quantum advantage, and careful resource estimates are continually revised as both hardware and algorithms improve. These advances guide the community's expectations and influence the design of next-generation quantum computers. For example, knowing that 1 million qubits might break RSA in a week [14] intensifies efforts to develop post-quantum cryptography now, and conversely assures stakeholders that the quantum hardware goal (though daunting) has a concrete pay-off once reached.

II. Physics of Qubit Implementation

Quantum computing hardware relies on diverse physical phenomena to realize qubits and quantum gates. Here we explain the physics behind four prominent implementations: superconducting Josephson junction qubits, trapped-ion qubits, neutral atoms in optical tweezers, and photonic qubits. Each approach uses very different physical systems to obtain two-level quantum states that can be controlled and entangled.

A. Josephson Junctions and Superconducting Circuits

Superconducting qubits are built from circuits made of superconducting materials (metals with zero electrical resistance at cryogenic temperatures). The key nonlinear element enabling qubit behavior is the **Josephson junction** – a sandwich of two superconductors separated by a very thin insulating barrier [20]. In a Josephson junction, Cooper pairs of electrons (pairs bound in the superconducting state) can tunnel quantum mechanically through the insulator [20]. This tunneling current has special properties: it can flow without any voltage (the Josephson effect), and it introduces a nonlinear energy relation in the circuit. In essence, the Josephson junction behaves like a nonlinear inductor with no energy dissipation [20]. This nonlinearity is crucial: it allows the circuit to have discrete quantum energy levels that are not equally spaced, so that one can isolate a two-level qubit subspace.

A common design is the **transmon qubit**, which consists of a Josephson junction in parallel with a capacitor. The junction's nonlinear inductance and the capacitor form an anharmonic oscillator. The lowest two energy levels of this circuit are used as $|0\rangle$ and $|1\rangle$ qubit states, while higher levels are off-resonant (due to the anharmonicity) [20]. By shunting the junction with a large capacitance, transmons are made insensitive to charge fluctuations (hence the name trans-mon for transmission-line shunted plasma oscillation qubit), dramatically improving coherence [20]. Another variant is the **flux qubit**, a superconducting loop interrupted by one or more Josephson junctions. In a flux qubit, the two qubit states correspond to current circulating clockwise vs counterclockwise in the loop (related to different quantized magnetic flux through the loop) [20]. The Josephson junctions in the loop create a double-well potential for the flux, enabling a two-level system that can be controlled by magnetic fields.

Physically, the qubit's state is represented by the phase difference of the superconducting wavefunction across the Josephson junction (for charge/parity qubits) or the quantized magnetic flux in a loop (for flux qubits) [20]. The system is operated at millikelvin temperatures in a dilution refrigerator to eliminate thermal excitations (transitions between qubit states due to thermal energy). Control is achieved by nanosecond-scale microwave pulses delivered to the circuit via on-chip antennas or transmission lines. A resonant microwave pulse at the qubit's transition frequency induces Rabi oscillations between $|0\rangle$ and $|1\rangle$, performing single-qubit rotations. Two-qubit gates are typically done by coupling qubits via a capacitor or mutual inductance; a controlled interaction (often using a tunable coupler or a bus resonator) entangles the qubits. For instance, a CZ or

iSWAP gate can be implemented by bringing two qubits into resonance or using parametric modulation of their coupling.

Measurement of superconducting qubits is done through **dispersive readout**: the qubit is coupled to a microwave resonator, and the qubit's state slightly shifts the resonator frequency. By sending a microwave probe and detecting its phase and amplitude, one can infer the qubit state ($|0\rangle$ vs. $|1\rangle$) from the resonance response – effectively a quantum non-demolition measurement.

The Josephson junction's role is pivotal – it provides a "nonlinear oscillator" that can be quantized. Unlike a normal LC circuit which has evenly spaced energy levels, a Josephson LC circuit has a level spacing difference that allows selecting only two levels to function as a qubit [20]. The junction permits superconducting phase to behave as a quantum variable. When two superconductors are separated by an insulator, a supercurrent can tunnel across with a well-defined relation $I = I_c \sin\phi$ (where ϕ is the phase difference). This leads to a Josephson energy $E_J(1 - \cos\phi)$, which provides the anharmonic potential for the phase "particle." The small capacitance of the junction provides a charging energy term. Together, the circuit's Hamiltonian yields quantized bound states for the phase difference or charge number – these quantized levels form the qubit basis [20].

In summary, superconducting qubits harness collective electron behavior (Cooper pairs) in a macroscopic circuit. The Josephson junction is the element that "**enables the formation of quantum states**" by allowing controlled tunneling of Cooper pairs and introducing nonlinearity [20]. The result is an engineered two-level quantum system that can be fabricated with lithography. The physics is an interplay of superconductivity (for coherence and zero resistance), quantum tunneling (in the junction), and microwave engineering (for control and readout). As a testament to their quantum nature, transmon qubits have shown quantum coherence times in the millisecond range and can maintain superposition and entanglement as needed for algorithms [3]. Continued improvements in materials (e.g., reducing dielectric loss, using 3D cavities, surface treating) and junction design aim to further improve coherence [20]. Superconducting qubits were among the first to demonstrate a range of quantum algorithms and remain at the forefront of building larger processors due to their fast gate speeds and compatibility with electronic integrated control.

B. Trapped-Ion Physics

In trapped-ion quantum computers, the qubits are physical ions – atoms that have been ionized to carry a net charge. Common examples are Ytterbium-171 or Beryllium-9 ions. The basic physics is that ions can be confined in space using electromagnetic fields. A linear **Paul trap** (RF trap) creates oscillating electric fields that produce a restoring force on the ions, effectively trapping them in a harmonic potential along a line. Multiple ions trapped in this potential naturally arrange in a linear crystal (due to mutual Coulomb repulsion balancing the trap forces). The ions are held in vacuum (ultra-high vacuum chamber) to avoid collisions with air molecules [8]. They are laser-cooled to very low motional energy – typically using Doppler cooling and resolved sideband cooling to reach the motional ground state [8].

Each ion provides an internal two-level system to serve as a qubit. A popular choice is a pair of hyperfine ground states of the ion (for example, in Yb^+ one can use two hyperfine levels separated by a microwave frequency – these are extremely stable "clock" states). Another choice is an optical transition: the ground state and a metastable excited state (e.g. $S_{1/2}$ and $D_{5/2}$ in Ca^+) separated by an optical frequency [8]. Hyperfine qubits have the advantage of long coherence (since both states are low-energy levels relatively insensitive to environment), whereas optical qubits allow fast optical transitions but have finite excited state lifetimes [8]. In any case, the qubit states are two distinct electronic energy levels of the ion.

How gates are done: Single-qubit gates on ions are performed by shining a laser or applying microwaves resonant with the qubit transition. A resonant laser pulse (with appropriate phase, frequency, and duration) will drive Rabi oscillations between the two levels, allowing arbitrary single-qubit rotations [8]. Because ions are well separated (typically a few micrometers apart in the trap), one can address individual ions by tightly focused beams or by frequency selection if using magnetic field gradients (each ion can be slightly detuned by a gradient, allowing frequency addressing) [8]. This gives precise single-qubit control with errors as low as 10^{-4} in experiments.

For **two-qubit entangling gates**, trapped ions exploit their collective motional modes. The ions in a trap vibrate together – for N ions, there are N collective modes of vibration (common center-of-mass mode, breathing modes, etc.). By using a two-qubit gate laser pulse that slightly off-resonantly excites a shared motional mode, one can induce an effective interaction between the internal states of two ions. The canonical two-qubit gate is the Mølmer-Sørensen (MS) gate: two laser beams create a state-dependent force on the ions. If both ions are in certain states, the force drives motion; the interference of forces entangles the spin states with motion transiently and then recombines them, resulting in an entangled spin state (e.g. $|00\rangle + |11\rangle$) independent of motion at the end. In simple terms, the MS gate uses a bichromatic laser field that causes a spin-flip conditioned on the vibrational excitation, entangling the qubits through the shared phonon bus.

Another common gate is the Cirac-Zoller controlled-NOT gate, which uses a sequence of laser pulses to map the qubit state of one ion onto phonon excitation and then onto another ion, achieving a CNOT. Both methods have achieved $>99\%$ fidelity in labs.

The physics of these gates relies on the fact that **any pair of ions can interact via the collective modes**, giving all-to-all connectivity in a small chain [8]. This is a stark contrast to many solid-state qubits. The Coulomb force coupling is long-range (all ions are coupled by their mutual repulsion), and by spectrally addressing different motional modes or by physically moving ions, one can perform entangling operations between arbitrary qubit pairs, not just nearest neighbors [8].

Readout of ion qubits is typically done by fluorescence detection. One of the qubit states (say $|0\rangle$) is chosen as the "bright" state that will emit light when a resonant laser is applied, while the other state $|1\rangle$ is a "dark" state (either it is not resonant with the light or it is a state that does not fluoresce). For example, in Yb^+, shining a laser resonant with a cycling transition from one hyperfine level causes photons to scatter if the ion is in that state, but will not if the ion is in the other hyperfine state. A photomultiplier or camera collects the fluorescence from each ion: bright = $|0\rangle$, dark = $|1\rangle$ [8]. By detecting many photons (or lack thereof) from each ion, high-fidelity projective measurement is achieved (with readout errors $\sim 0.1\%$ in the best systems).

Trapped-ion qubits are extremely pristine quantum systems – each ion is identical (being the same atomic species), eliminating fabrication variation [8]. Their coherence times can be superb; hyperfine qubits have T_2 coherence times of tens of minutes or even hours in magnetic-field-stabilized traps, limited mostly by magnetic field noise (which can be suppressed using "clock" states that are field-independent) [8]. The primary decoherence mechanisms in ion traps are magnetic field fluctuations (for Zeeman qubits), photon scattering from lasers (for optical qubits during illumination), and motional decoherence (heating of the trap). But experimentally, ions have shown some of the longest coherence of any qubit.

Physics summary: Trapping uses dynamic electric fields to confine ions; cooling uses laser-photon momentum kicks to remove kinetic energy; qubit states are internal electronic states (manipulated by laser photons causing transitions); entanglement arises by coupling internal states to shared vibrational quanta (phonons) via laser forces. All operations happen with the ions suspended in free space (no touching surfaces), which contributes to their low error rates. The cost is that operations are relatively slow (gates in microseconds, and one often needs to wait for motion to settle between gates, although parallel operations on different pairs are possible

to an extent).

Trapped ions have also realized small-scale quantum error correction (e.g., a 7-ion logical qubit using the [[7,1,3]] code where a logical memory was improved by QEC) [8], and multi-ion GHZ states up to 14 ions have been created as a demonstration of entanglement generation. As of 2024, academic and commercial ion traps routinely juggle up to $\sim$20 qubits with high fidelity, and the focus is on modular scaling (networking multiple traps). The QCCD approach may move ions between zones for scaling [8], and indeed Honeywell (Quantinuum) has realized multi-zone traps where ions are shuttled through junctions on chip. Another approach is to use photonic interfaces: each ion trap can emit photons entangled with its ions, and interfering those photons between traps can entangle ions in different traps (remote entanglement), effectively creating a scalable network of ion qubits [8].

In summary, the trapped-ion implementation leverages well-understood atomic physics: electromagnetic trapping of charged particles and laser spectroscopy for control. The qubit is literally an atomic clock transitioned into the quantum computing paradigm. This gives it excellent reliability and predictability – indeed, trapped ions have naturally identical qubits and high-fidelity operations as one of their strongest selling points [8].

C. Neutral-Atom Optical Tweezer Qubits

Neutral atom qubit systems share some similarities with trapped ions (both use atoms and laser control), but differ in the method of trapping and entangling. Instead of using charge to trap, neutral atoms are trapped by optical forces. A highly focused laser beam can create a dipole trap: the electric field of the laser induces a dipole moment in the atom, and if the laser is red-detuned (frequency below the atomic resonance), the atom is drawn to the region of highest intensity (the focus). These traps are nicknamed "optical tweezers." By steering multiple laser beams (using acousto-optic deflectors or spatial light modulators), one can trap hundreds of neutral atoms in arbitrary configurations (1D, 2D arrays, or even 3D) [9].

The typical procedure begins with collecting atoms in a magneto-optical trap and cooling them (similar to ions, Doppler cooling is applied) [9]. Then the atoms are loaded into an array of optical tweezers – often initially some traps are empty (occupancy is random). A fluorescence image is taken to see which traps are filled [9]. Then atoms are rearranged: using optical tweezers as "optical conveyor belts," atoms can be moved so that a chosen set of traps are filled with exactly one atom each [9]. This yields a register of qubits (each trap = one qubit).

For qubit states, neutral atoms can use hyperfine ground states (like $|0\rangle$ and $|1\rangle$ are two hyperfine levels of the ground state, similar to ion hyperfine qubits) or one ground state and one long-lived excited state. A particularly important aspect is the use of **Rydberg states** for entanglement. A Rydberg state is an excited state of an atom with the valence electron promoted to a very high principal quantum number (e.g., n $\sim$ 50-100). In such states, the electron is far from the nucleus, and the atom has a huge electric dipole moment. Two Rydberg atoms can interact strongly over distances of several micrometers via dipole-dipole or van der Waals interactions. This leads to the **Rydberg blockade** phenomenon: if one atom is excited to a Rydberg state, any other atom within a certain blockade radius (perhaps 5-10 μm) is shifted out of resonance and cannot be excited to the Rydberg state [9]. By exploiting this, a controlled two-qubit gate can be realized: a laser tries to simultaneously excite two atoms to Rydberg state; the blockade ensures that at most one gets excited. With a proper sequence (excite control atom to Rydberg, conditional excitation of target, then de-excite control), one can implement a CNOT or CZ gate depending on pulse sequences. Essentially, the presence or absence of a Rydberg excitation on one atom controls the transition of another – that entangles their qubit states [9].

Single-qubit gates on neutral atoms are done with microwaves or two-photon Raman transitions between the

two ground states, or direct optical transitions if using one excited state as $|1\rangle$. These are similar to ion single-qubit operations (Rabi rotations driven by resonant fields).

One significant benefit of neutral atoms is **geometric flexibility**: atoms can be arranged in 2D layouts which can mirror the structure of error-correcting codes (like a 2D grid for the surface code) or problem graphs (for quantum optimization algorithms) [9]. Connectivity can be reprogrammed by moving atoms or by using Rydberg interactions that can reach neighbors a few sites away (longer-range connectivity than fixed nearest-neighbor). Neutral atoms also allow mid-circuit reconfiguration – one could move qubits during computation potentially, or swap atoms, since they are mobile to an extent within the optical traps.

Readout for neutral atoms is similar to ions: state-dependent fluorescence. Usually, one hyperfine state is coupled to a cycling transition and fluoresces under laser illumination, while the other state remains dark [9]. A camera or photodetector array can read out entire 2D arrays in one shot by imaging fluorescence spots.

Decoherence in neutral atoms can be very low. Ground-state hyperfine coherence can be on the order of seconds (limited by magnetic field noise or laser phase noise for the Raman transitions) [9]. Rydberg states, being highly excited, have much shorter lifetimes ($\sim$100 microseconds or so due to spontaneous decay), but they are only occupied briefly during gates. Achieving high two-qubit gate fidelity in neutral atoms requires precise laser control to avoid phase errors and maintaining the blockade effectiveness. As of now $\sim$99% fidelity two-qubit gates have been reported [9], and there is active work to push this higher through better laser stability, cooling the atoms to lower motional states (to reduce thermal motion causing phase shifts), and advanced pulse shaping.

Neutral atom arrays have already been used for quantum simulation of many-body physics (e.g., simulating Ising magnetism by turning on controlled Rydberg interactions across an array). They have also been used to solve small instances of optimization problems like maximum independent set by encoding the problem graph into a Rydberg interaction graph [9] – the Rydberg blockade natively enforces that two nearby atoms cannot both be excited, which maps to constraints in combinatorial problems [9]. These demonstrations show the versatility of neutral atom qubits for both digital and analog quantum computing.

Physics summary: Neutral-atom qubits physics involves laser trapping (optical dipole forces), which is a purely optical means of confinement as opposed to electromagnetic fields in ion traps. The qubit manipulation is largely through laser-driven transitions and the exploitation of extremely high-lying atomic states (Rydberg states) to generate interactions. The appeal of this method lies in its potential scalability (hundreds of traps can be created with advanced optics) and its naturally long qubit coherence. The complexity lies in controlling many laser beams in parallel and dealing with atomic motion (atoms are not fixed rigidly; even at micro-Kelvin they have slight motion within traps, requiring robust gate schemes). The field has advanced to the point where companies plan 1000-qubit neutral atom systems within a couple of years [9], which would make it one of the highest-qubit count platforms, though error rates will need to keep improving to fully use those qubits in algorithms.

D. Photonic Qubit Generation and Detection

Photonic qubits are fundamentally different from matter-based qubits: they are quanta of light, which means they are always traveling at the speed of light (unless stored in a delay line or optical cavity) and do not interact with each other in linear media. The physics of photonic quantum computing centers on generating single photons in well-defined quantum states, interfering photons to create entanglement, and reliably detecting photons to perform measurements.

Single-Photon Sources: An ideal single-photon source emits one and only one photon on demand, with high purity and indistinguishability (so that multiple photons can interfere). Common approaches:

- **Spontaneous Parametric Down-Conversion (SPDC)**: A nonlinear crystal is pumped with a laser; occasionally, one pump photon splits into two lower-energy photons (signal and idler) in an entangled pair. By heralding (detecting one photon of the pair), the other is prepared in a single-photon state. SPDC sources are widely used in experiments, but they are probabilistic (each pump pulse has a small chance of producing a pair, and occasionally produces 0 or 2+ pairs, leading to trade-offs).
- **Quantum Dots**: Semiconductor quantum dots can act as "artificial atoms" that emit single photons. When a quantum dot is excited, it can relax by emitting a single photon. Placing a quantum dot in an optical microcavity can enhance emission and funnel photons into a single mode for high efficiency. Recently, quantum dot sources have shown near-deterministic single-photon emission with high purity and indistinguishability (fidelities $> 99\%$).
- **Other sources**: Single atoms, defects like NV centers in diamond, or molecules can emit single photons on demand, but coupling them to optical fibers or circuits is challenging. Silicon photonic chips with integrated ring resonators have also been used to generate photon pairs via four-wave mixing (another nonlinear optical process).

PsiQuantum and others are pursuing large numbers of on-chip SPDC or quantum dot sources to generate the thousands or millions of photons needed for a cluster-state machine [21]. Quandela's photonic platform, for example, uses quantum dot sources integrated on chip [21].

Entanglement and Two-Qubit Gates: Photons do not directly interact (two photons can pass through each other without disturbance). However, if two photons are indistinguishable and meet at a beam splitter, they undergo Hong-Ou-Mandel interference – a quantum effect where they bunch into the same output path in an entangled way. Using additional ancilla photons and post-selection (measuring certain outcomes), one can implement probabilistic two-qubit gates like CNOT (as in the Knill-Laflamme-Milburn (KLM) scheme). These gates only succeed some fraction of the time, but if you have offline resource states (entangled photons) you can try until you succeed. This is resource-intensive but in principle scales.

A more practical approach for large scale is **measurement-based quantum computing (MBQC)**: instead of performing gates by two-photon interactions, one prepares a large entangled state of many photons (a cluster or graph state, e.g. a 2D lattice of entangled photons) offline. Then, quantum computation is carried out by performing single-photon measurements on this cluster state in a certain sequence, which effectively enact quantum gate operations on the unmeasured part of the cluster. The cluster state serves as a substrate for computation, and measurements (with classical feed-forward based on outcomes) drive the logic of the algorithm. Generating a massive cluster state is the main challenge – it requires entangling many photons. One approach is to fuse smaller entangled states into a big one using linear optics and detection (fusion gates). PsiQuantum's blueprint involves creating small (e.g. 5-photon) entangled resource states on chip and then fusing them with a network of interferometers to build a large 3D cluster state suitable for fault-tolerant MBQC [22].

The physics of generating entanglement thus often reduces to interfering photons on beam splitters and using **detective measurement to induce effective nonlinearities**. Because direct nonlinear interactions at the single-photon level are very weak in most media, using measurement to induce entanglement has been the dominant method.

Detection of Photonic Qubits: Measuring a photonic qubit typically means detecting the presence or polarization/path of a photon. Single-photon detectors are a critical component. Leading detectors are:

- **Superconducting Nanowire Single-Photon Detectors (SNSPDs)**: A superconducting wire biased near its critical current will momentarily go normal (produce a voltage pulse) when even a single photon is absorbed and deposits energy. SNSPDs can have detection efficiencies $> 90\%$ and very low dark count rates, with timing resolution of tens of picoseconds. They usually operate at $\sim$2-3 K (liquid helium or closed-cycle cryocooler).
- **Avalanche Photodiodes (APDs)**: Semiconductor diodes biased above breakdown that produce a large current pulse when a photon creates an electron-hole pair. APDs are often used at visible or telecom wavelengths, though their efficiency and noise are not as good as SNSPDs.
- **Transition Edge Sensors (TES)**: Superconducting calorimetric detectors that can even number-resolve (tell how many photons arrived by measuring the precise heat).

High-efficiency detection is crucial because lost photons = lost qubits. Photonic computation often requires detecting many photons during the computation (for feed-forward). If each detection has, say, 95% efficiency, across many steps the success probability multiplies, so pushing detection to $>$99% efficiency is an engineering goal.

In practice, current photonic experiments often require cryogenic setups for detectors (even if the photonic circuits themselves are at room temperature). Thus, a full photonic quantum computer might still have significant cryogenic infrastructure – for example, arrays of SNSPDs integrated near the photonic chip.

Physics summary: In photonic qubit physics you create photons one by one (ensuring they are indistinguishable and in well-defined modes), you route them through a network of beam splitters and phase shifters to create entanglement (either during gates or upfront as cluster states), and you measure them with single-photon-sensitive detectors. Photons are robust carriers of quantum information (low decoherence), which is advantageous for connecting distant modules (e.g. distributing entanglement over fibers) and for room-temperature operation. The difficulty lies in the probabilistic nature of entanglement operations and the need for a large overhead of photons (many more photons typically need to be generated than the number of logical qubits, because many get used/lost in the creation of entanglement).

Recent achievements include a photonic 12-qubit entangled state used to demonstrate a small quantum algorithm, Gaussian boson sampling experiments with 50-100 photons (as quantum advantage demonstrations), and the integration of photonic components on silicon chips for stability and scalability [21]. **In 2025, PsiQuantum announced they are working toward a million-qubit photonic quantum computer using silicon photonics and have begun building a facility** for it [16]. This involves thousands of photon sources and detectors integrated together, something made possible by leveraging semiconductor fabrication techniques.

In closing, photonic qubits represent a fusion of quantum optics and optical engineering. The physics is grounded in quantum optics principles established in the 20th century (like interference and entanglement of photons), now pushed to extreme scales. If successful, photonic quantum computers might look more like optical communication systems – with lasers, waveguides, and detectors – than traditional computers, **and they could operate at room temperature** aside from certain components. They are also natural for connecting quantum machines, since photons are the information carriers in quantum networks (e.g., quantum internet proposals count on photons to link nodes).

III. Systems Engineering for Scaling to One Million Qubits

Building a quantum computer with one million physical qubits is an enormous engineering challenge. It is not just a matter of qubit counts; it requires careful design of power and cooling systems, control electronics, error correction infrastructure, interconnects between qubit modules, and even consideration of the building

and personnel needed. In this section, we discuss the systems-level issues: cryogenic and power requirements, classical control electronics and wiring, the overhead of error correction and networking qubits, and facility-scale considerations including energy sources like gas vs. nuclear power for such a large machine.

A. Power and Cryogenic Requirements

Cryogenics: Most leading qubit platforms (superconducting qubits, spin qubits, some photonic detectors) require cryogenic operation at very low temperatures. Superconducting qubits operate at ∼10-20 millikelvin, which is achieved by **dilution refrigerators**. A million-qubit superconducting quantum computer likely cannot reside on a single chip in one refrigerator due to size and wiring limits; it would be partitioned into many modules each in its own cryostat, or one extremely large cryostat housing multiple chips. IBM has already developed a concept for a "super-fridge" named Goldeneye – a dilution refrigerator 10 feet tall and 6 feet wide with an experimental volume of 1.7 cubic meters, intended to accommodate large quantum systems [17]. **Goldeneye** can hold multiple refrigeration units and is designed to support cooling power on the order of 10 mW at 100 mK (and >20 W at 4 K) – substantially more than standard lab fridges [17]. This demonstrates the engineering direction for scaling: building larger or networked cryogenic systems to house many qubits. A million qubits might be organized as, for example, 1000 modules of 1000 qubits each, requiring perhaps hundreds of dilution refrigerators.

Each dilution refrigerator itself is a complex machine consuming significant power (for vacuum pumps and helium compressors). Today's large cryostats might consume on the order of ∼20-50 kW of electrical power each. If hundreds of such units are needed, the cryogenic cooling alone could draw **tens of megawatts**. Alternatively, new cooling technologies might be developed: for instance, closed-cycle cryocoolers in a centralized **helium liquefaction plant**, or novel cryogen-free dilution units (like the Bluefors Kide platform IBM mentions for its System Two) [17].

Additionally, trapped-ion and neutral atom systems, while often operated at or near room temperature for the atoms themselves, may require certain components at low temperature (superconducting detectors for photons, ultra-stable laser reference cavities at cryo, etc.). But compared to superconducting qubits, ion/atom platforms can potentially avoid large cryogenic loads for the qubits, instead trading for lasers (which have their own overhead).

Power Requirements: The total power for a million-qubit quantum computer includes:

- **Cryogenics (if needed)**: as described, likely in the megawatt range if using dilution refrigerators for many modules.
- **Control electronics**: Generating microwave pulses, laser beams, or bias currents for one million qubits will be power-intensive. Currently, a single superconducting qubit's control line might use microwatts in the qubit but much more in room-temperature electronics (the amplifiers, AWGs, etc.). For a million qubits, control electronics must be scaled and likely moved closer to the qubits (cryo-CMOS at 4 K) to avoid an impractical heat load from wiring (one million coaxial cables is impossible to accommodate).
- **Wiring and Filtering**: Each qubit typically has one or more control lines and a readout line. In current devices, each coax line into the fridge carries a heat load of a few milliwatts at mK stage (after filtering, attenuators, etc.). One million qubits clearly cannot have one million separate coax cables; new approaches are needed such as multiplexing many qubits per line (RF multiplexing), photonic fiber feedthroughs, or in-situ control chips. Nonetheless, whatever wiring scheme is used, the control system will dissipate power. Cryogenic control chips (CMOS at low temperature) have to operate within the cooling power budget of the fridge, which might be only a few milliwatts at the coldest stage. This implies that much of the control

must be at higher temperature stages (e.g., 4 K or 40 K stages with a few watts available). For one million qubits, advanced integrated control electronics are mandatory – as Google's Shirin Montazeri emphasized, "achieving large-scale quantum computation necessitates control electronics capable of programming millions of qubits with high fidelity", representing a significant engineering challenge and opportunity [3].

Given these factors, **the power consumption of a million-qubit system could rival that of a small classical supercomputer data center**. We might estimate **on the order of 10-20 MW** total, though this is speculative. For comparison, the current top classical supercomputers (exascale machines) consume ∼20-30 MW of power. A quantum computer might eventually be more energy-efficient for certain computations, but the overhead of cooling and control erodes some of the advantage. Still, if a quantum computer can solve a decades-long classical computation in hours, the energy per task could be far lower even if the power draw is high.

Heat and Cooling Infrastructure: Removing heat from tens of megawatts of equipment requires robust HVAC (for room-temperature equipment) and possibly liquid cooling loops for high-density electronics racks. The cryogenic helium compressors often require chilled water. In essence, a million-qubit quantum data center might look similar to an HPC data center with chillers, cooling towers, backup generators, etc., plus the unique dilution refrigeration infrastructure.

B. Control Electronics, Qubit Interconnects, and Error Correction Overhead

Control Electronics: The classical electronics that drive the quantum machine are critical. As devices scale, a key goal is to move from bulky laboratory instruments to integrated control systems. Already, companies are developing cryogenic CMOS chips that sit at 4 K to generate qubit control pulses (to drastically reduce the wiring bottleneck). For example, researchers have prototyped multiplexed control chips that can drive multiple qubits using one input line. Also, technology like FPGAs and DACs are being cryogenically qualified. The IBM System Two architecture envisions some control at 4 K and an internal cryogenic interconnect between multiple chips [17]. Google and other teams similarly are exploring superconducting resonator networks to distribute clock and control signals inside the fridge.

One big challenge is the precision required: controlling millions of qubits means billions of control parameters (pulse timings, amplitudes) which all must be calibrated and stable. Ensuring signal integrity (low noise, low crosstalk) when routing signals to so many qubits is nontrivial. This motivates advanced packaging like superconducting flex cables, on-chip filtering, etc.

Montazeri's note summarizes this: "A key challenge in realizing fault-tolerant QC lies in developing high-precision control electronics for qubit control and readout…achieving large-scale quantum computation necessitates control electronics capable of programming millions of qubits with high fidelity." [3]. It's an area where conventional RF engineering meets new cryo constraints, and will require innovations in low-power RF, high-density interconnects, and perhaps new materials (for example, using superconducting interconnects to carry signals with minimal loss into the cryostat).

Error Correction Overhead and Architecture: A million physical qubits will likely not be used to run a million-qubit algorithm directly; rather, they may encode on the order of a few hundred or a few thousand logical qubits under error correction. For instance, IBM's 2029 goal is ∼200 logical qubits (fault-tolerant) from a machine likely containing many thousands of physical qubits [13]. In a surface code, roughly ∼1000 physical qubits might encode 1 logical qubit with a target error rate, although exact numbers depend on physical error rates. Thus, the architecture likely involves a 2D array of qubits with modular tiling of error-correcting code patches.

The overhead of quantum error correction places demands on interconnects: Qubits need to interact with neighbors frequently for syndrome measurements. In a distributed multi-chip system, this means adjacent chips must be linked (for example, with superconducting bridges or via optical fiber if cryogenic photonic interconnects are used). **Modular architectures** are being explored: instead of one monolithic chip, one could have many smaller chips each handling a region of the code, connected by high-fidelity links (maybe microwave waveguides or even using photons to connect distant logical qubits between cryostats).

Another aspect is **communication between logical qubits**. If one million qubits are all local, fine; but if they are split into 1000 modules of 1000 qubits, one needs inter-module entanglement distribution. This is reminiscent of a quantum network or a cluster of quantum processors. Technologies like photonic interconnects (where a superconducting qubit's state is mapped to a photon and sent via fiber to another cryostat to entangle with a distant qubit) are being developed. Even within one fridge, if chips are on different boards, one might use microwave-to-optical transducers to link them. Some efforts (e.g., at AWS and University of Chicago) have demonstrated **cryogenic microwave links** between qubits in separate cryostats or through circulators, as a step toward a quantum intranet of sorts.

In ion traps, interconnect means photonic links between ion trap modules – which have been shown with a few modules entangled via fiber photons (IonQ and Duke University experiments). For neutral atoms, it could mean physically transporting atoms from one array to another, or using photons as well.

Error correction also requires fast classical processing: extracting error syndromes from many measurements and feeding back to correct errors in real-time. In a surface code, for example, each round involves measuring parity checks and then doing a lookup or decoding algorithm. Recent work increasingly treats *decoder latency* as a first-class constraint, motivating hardware-friendly approaches (e.g., belief-propagation variants such as Relay-BP for qLDPC decoding) and demonstrations that decoding can be kept on the microsecond timescale using dedicated classical accelerators [28]. For a million qubits, millions of syndrome bits per second (or more) will be produced. Classical co-processors (FPGAs, GPUs, or specialized decoders) need to crunch this data with low latency (microseconds timescale) to send correction signals (like flipping a qubit or adjusting a gate sequence). This classical processing might be distributed across the machine to avoid bottlenecks.

Thus, the system architecture becomes a **heterogeneous computing system**: a quantum processor networked with a fast classical control plane.

Interconnect Bandwidth: Consider that if qubits are producing readout signals at (say) 1 kHz (typical for QEC cycles), a million qubits produce 1e6 measurements per millisecond, i.e., 1e9 per second. Handling this data – both transmitting it out of the cryo and processing it – is a formidable task. It implies very high-bandwidth connections from the cryostat to classical computing resources. This is one reason to push some processing inside the fridge (at 4 K or even 40 K stages) so that not all raw data leaves the cold environment – only aggregated decisions.

In short, scaling to a million qubits **forces a re-think of the entire stack** from qubit chip to outside world. Techniques like **multiplexing, on-chip control, 3D integration, and quantum-classical co-design** are active research areas. Fortunately, early prototypes of these ideas exist (e.g., Imec demonstrated a 300 mm wafer process for integrating spin qubits with cryo-CMOS control lines in proximity [3]).

C. Facility Infrastructure: Footprint, Staffing, and Real Estate

Constructing a facility for a million-qubit quantum computer is akin to a **major scientific apparatus or data center project**. It involves not just the machine itself but also supporting infrastructure and human expertise.

Facility Contractors: Given the scale, sensitivity, and engineering demands of a million-qubit quantum computing facility, **Jacobs** and **Parsons** emerge as strong candidates to serve as prime contractors for the construction of both the quantum processing centers and the supporting high-performance computing (HPC) and data infrastructure. Both firms have deep experience executing mega-scale, mission-critical government computing projects. Jacobs, for example, has delivered secure, energy-intensive facilities for the Department of Energy (DOE), Department of Defense (DoD), and Intelligence Community (IC), including design and operations support for supercomputing installations at national laboratories such as Oak Ridge and Lawrence Livermore. Similarly, Parsons has a proven track record with cyber-secure defense infrastructure and recently expanded into advanced computing systems integration with multi-billion-dollar programs under the NSA and military branches. Both companies have made public statements signaling their strategic interest in quantum readiness, emphasizing capabilities in high-reliability electrical systems, advanced cryogenic plant engineering, and modular, vibration-isolated cleanroom environments – key components for housing quantum processors that operate at millikelvin temperatures. Their portfolios include facilities with 100+ MW power footprints, precision environmental controls, and scalable data integration with classical HPC clusters – all crucial for co-locating error-corrected quantum modules and the real-time decoding hardware required to support them. As the U.S. seeks to realize national-scale quantum infrastructure, Jacobs and Parsons bring the specialized experience, clearances, and industrial capacity required to deliver secure, scalable, and technically advanced quantum data centers.

Space Requirements: Quantum computing hardware, especially when you include multiple cryostats, racks of electronics, and cooling systems, is not small. For instance, IBM's current Quantum Computation Center in Poughkeepsie houses numerous dilution refrigerators (each roughly closet-sized) in a room with extensive cabling and control racks. A million-qubit system might occupy an entire building or a large laboratory floor. As a rough guess, if one dilution fridge (like Goldeneye) is $\sim$3 meters tall and $\sim$2 m diameter, and maybe a dozen of those are needed, that's already tens of square meters just for fridges, plus clearance and support equipment. PsiQuantum's planned facility in Illinois (the Illinois Quantum and Microelectronics Park in Chicago) is expected to be a sizable site – the press mentions a **43,000 ft^2 (4000 m^2) quantum computing laboratory** for a photonic quantum computer development [23]. This indicates thousands of square meters may be allocated for a single multi-million-qubit machine and its R&D space.

Environmental control is another factor. Vibration isolation might be needed for sensitive equipment (ion traps and atom arrays are sensitive to vibrations; superconducting qubits less so, but still, mechanical stability of cryostats is considered). Temperature and humidity control for laser labs, electromagnetic shielding for low-noise environments, and cleanroom-like conditions for some parts (to avoid dust affecting optics or qubits) could be needed. All this influences the facility design.

Workforce and Expertise: Operating a million-qubit quantum computer would require a multidisciplinary team. For comparison, large particle accelerators or fusion experiments involve hundreds of PhDs and engineers for operation. A quantum computer may not need that many once it's stable, but at least dozens of highly skilled personnel: quantum physicists to calibrate qubits and ensure performance; microwave engineers; cryogenic engineers to maintain and troubleshoot dilution refrigerators; classical computing engineers for control systems; software developers for higher-level programming, etc.

PsiQuantum's partnership with Illinois expects to create "at least 150 jobs in the next five years" for their

initial operations [16], including roles from PhD scientists to technicians [16]. This gives a sense of scale – likely several hundred staff in steady state to run a large quantum computing facility, similar to a big data center or research lab.

Moreover, continuous maintenance will be needed. Qubits need calibration (currently daily or more often calibrations are standard for a few-qubit devices; one hopes automation and improved stability will reduce this, but with a million qubits, even rare drift could mean something is always being calibrated). Automated calibration routines running on classical computers will likely handle much of it, but human oversight remains important.

Supporting Infrastructure: The facility will need robust electrical power supply (discussed below in Section III-D), backup power (to safely maintain cryogenic systems in power outages or to at least warm up safely – note that a sudden loss of power in a dilution fridge can cause rapid boil-off of helium, etc., so backup generators or uninterruptible power might be used). It may also need high-capacity data connections if the quantum computer is to be cloud-accessible (users sending jobs to it), meaning fiber optic links and significant classical compute integration.

If using complex equipment like high-power lasers (for ions/atoms), one needs proper cooling for those laser systems and possibly UV-resistant lab enclosures. If using high-pressure cryogens or helium handling, safety systems (oxygen deficiency monitors in case of helium leak) must be in place.

Real Estate Location: Preferred site characteristics might include:

- Proximity to a tech talent pool (quantum experts, engineers). Many quantum centers are near universities or tech hubs.
- Stable seismic and geological conditions (to reduce vibration and risk of earthquakes; one wouldn't site a delicate quantum lab on a fault line if avoidable).
- Reliable power grid and possibly cool climate (to reduce cooling costs). Some data centers choose cooler climates to ease HVAC loads – similarly, a quantum facility might benefit from a location where chillers operate more efficiently.
- Political and economic support: e.g., Illinois offered partnership to PsiQuantum, presumably including incentives. Areas with existing advanced research infrastructure (like national labs or semiconductor fabs) are attractive. For example, IBM's Poughkeepsie site is an IBM campus with needed utilities; Google's Santa Barbara lab is near a university with a history in quantum; national labs (like Oak Ridge, Argonne) might host future large systems as well.

Quantum computing is also an international endeavor, so there might be considerations of security (if it's breaking encryption, perhaps a secure site). The U.S. might prefer domestic sites for major quantum facilities.

In essence, building a million-qubit computer could be comparable to projects like a new supercomputer center or even a large physics experiment facility. It needs significant planning in construction, with reinforced floors for heavy equipment, possibly additional electromagnetic shielding in walls (to reduce external EM noise), and so on.

D. Power Source Considerations (Gas vs. Nuclear)

Given the projected power consumption (potentially tens of megawatts) and the critical need for uninterrupted operation (quantum experiments can be spoiled by sudden shutdowns, and cryogenic systems are best kept running continuously), the choice of power source for a quantum computing facility is important. Two options mentioned are on-site gas generation and on-site nuclear, in addition to drawing from the grid.

- **Grid Electricity**: The simplest is to use the local grid, possibly with redundancy (two separate feeds from different substations, etc.) plus uninterrupted power supply (UPS) and diesel generators for backup. This is how most datacenters operate. If the quantum facility is near, say, a large city or an industrial park, grid power is feasible. However, a quantum data center might want higher reliability than typical – a momentary outage could disrupt experiments and cause cryostats to begin warming. So some local generation or energy storage might be used.
- **On-site Gas Turbines**: A natural gas generator (or combined heat and power plant) on-site could provide a steady power supply and reduce dependence on the grid. Gas turbines can be quite efficient especially if waste heat is reused (for building heating or absorption chilling). Gas infrastructure is common and fuel is generally available in most places. The downside is carbon emissions and fuel price volatility. But from a pure technical standpoint, many large facilities (hospitals, universities) have their own gas cogeneration for reliability. A quantum center could similarly use gas to ensure consistent power and maybe even use the exhaust heat for something (though quantum cryogenics can't directly use waste heat, it could help with building climate control, etc.).
- **Nuclear (Small Modular Reactor or Microreactor)**: A nuclear source on-site is a more exotic choice but is gaining attention for datacenters and research facilities aiming for zero-carbon, reliable power. A small modular reactor (SMR) in the tens of MW range could theoretically supply a quantum computing facility with stable power for years without refueling. The benefits are high reliability (nuclear plants have very stable output and do not depend on weather or frequent refuel) and zero carbon emissions at point of generation. However, nuclear has substantial regulatory and safety burdens. Siting a nuclear reactor (even a small one) requires approval and likely a remote or well-secured location. It might only be practical if the quantum facility is part of a larger national lab or in an area with a nuclear research allowance. The former US Steel South Works site in Chicago (for PsiQuantum's facility) is not nuclear powered; it likely will use grid power with upgrades.

Nuclear could make sense if quantum computers become critical national infrastructure requiring guaranteed uptime and minimal environmental impact. For example, if by 2035 a million-qubit quantum computer is an invaluable asset, one might treat it somewhat like how some supercomputers are considered critical – though they typically don't go as far as dedicated power plants; the xAI datacenter in Tennessee is a recent exception.

One idea is pairing a quantum center with a **renewable energy plus battery** installation for sustainable power. But since quantum machines will likely run continuously (not just when the sun shines or wind blows), backups or grid still needed. Nuclear provides base load nicely.

Another factor: Large cryogenic systems might benefit from co-located **liquid helium plants** or advanced cryocoolers. Those also consume power. If nuclear, one might use some excess heat for things like driving absorption refrigerators or pre-cooling (though generally you need cold, not heat).

In conclusion, **gas vs. nuclear** is a trade-off of flexibility vs. long-term stability. Gas turbines can be installed relatively quickly and scaled (you can have multiple 5 MW generators for redundancy, etc.), and they can be turned on/off as needed. Nuclear is a long-term commitment but offers unparalleled reliability and no greenhouse emissions during operation. If the quantum computing industry pushes towards carbon-neutral operation (which could be a goal given many tech companies have sustainability pledges), nuclear or renewables would be considered over just burning gas. On the other hand, initial prototypes will probably just plug into the grid.

Example Scenarios:

- IBM's Poughkeepsie center likely uses grid power from NYSEG and has backup diesel generators (typical

datacenter).
- A future "Quantum SuperCenter" might use a dedicated SMR providing 50 MW, which could support multiple quantum machines and perhaps also classical supercomputers in the same campus (since classical HPC and quantum might complement each other).

Safety and Regulatory: If quantum facilities are near urban tech hubs (e.g., Chicago for PsiQuantum, Paris for Alice&Bob lab), nuclear on-site is unlikely due to regulatory restrictions in populated areas. Gas is more feasible anywhere there's a pipeline.

Cost considerations: Gas is cheaper initially, nuclear is expensive upfront but can be cheaper over long periods for fuel. If the facility expects to run for decades at high power, nuclear might pay off in stable electricity cost, whereas gas could fluctuate with market. Governments might subsidize a nuclear installation if the quantum facility is a strategic national asset.

Thus, the decision can depend on the specific use-case: a commercial quantum cloud provider might just use grid and offset carbon by buying renewables; a government quantum lab might integrate an SMR for independence.

In summary, scaling quantum computing is **not just a scientific challenge but a systems engineering challenge at the level of power and infrastructure**. These considerations illustrate why some have compared building a large quantum computer to historic large-scale projects like the space program or the Manhattan Project – **it will require coordination of many engineering disciplines and substantial resources**.

IV. Roadmap to a Million-Qubit Quantum Computer

Drawing together the current state-of-art and systems considerations, we can outline a roadmap of how to progress from today's devices (tens to a few hundred qubits) to the goal of one million qubits. This roadmap includes the technical milestones, expected timelines (as projected by major organizations), the human and capital resources needed at each stage, and even the desirable attributes of sites that will host these machines.

A. Technical Steps and Milestones

Near Term (2023-2025): Demonstration of Quantum Advantage and Small-Scale Error Correction. In this period, the focus is on reaching **error correction's dawn** – showing that logical qubits can outperform physical qubits. Already by 2023, Google showed a logical qubit with lifetime beyond any single physical qubit in the device [18]. We expect further demonstrations of logical operations (like a logical two-qubit gate between logical qubits) using $\sim$50-100 physical qubits per logical qubit in surface codes. Also in this timeframe, quantum advantage in specific domains may be demonstrated – e.g., a 100+ qubit quantum simulation that outperforms classical simulation, or a quantum machine learning task that a classical computer can't easily match. IBM and others are working toward demonstrating a provable quantum advantage in practical tasks by 2024-25 [1]. On hardware, IBM introduced the 433-qubit Osprey in 2022 and a 1121-qubit Condor chip in 2024 [3]. If achieved, that crosses the 1000-qubit mark in a single device. Techniques for scaling like 3D integration (multi-layer wiring) and modular coupling (entangling two chips together via chip-to-chip links) will appear. Indeed, IBM has introduced a modular approach in IBM Quantum System Two, integrating multiple processors within a shared cryogenic system and advancing cryogenic interconnect concepts for scaling [26].

By end of 2025, we expect:

- 1000+ physical qubit machines available (IBM, maybe Google).
- First generation **quantum interconnects** demonstrated (entangling qubits between cryostats or between distant modules).
- Topological-qubit pathway (high-variance): Microsoft announced the **Majorana 1** program and a prototype topological-processor roadmap aimed at very large scale [15]. However, repeatable, fault-tolerant *topological* qubits with full logical operations remain an active research goal; near-term milestones would be reproducible signatures, control primitives, and clear error-rate characterization before such devices can be credibly compared with leading superconducting, ion, or atom platforms.

Mid Term (2026-2030): Fault-Tolerant Prototype and Scaling to $\sim$1,000 Logical Qubits. This period is crucial for building a **useful, error-corrected quantum computer**. Google has publicly aimed for a practical error-corrected quantum computer by 2029 [12]. IBM similarly targets 2029 for a "practical" fault-tolerant machine (codenamed **IBM Quantum "Starling"**) with about 200 logical qubits [1]. So around 2028-2030, we might see the first **fault-tolerant prototypes** – these would be machines that incorporate full error correction on all qubits and can run long algorithms reliably. The size might be tens of thousands of physical qubits organized into a few hundred logical qubits.

To get there:

- Hardware scale-up will continue: IBM plans a 10,000+ qubit system by 2027-2029 [1] (some sources mention a 10k qubit goal by 2029). This likely involves multiple chips in a single cryostat or even a multi-cryostat network.
- New qubit types might come online. If Microsoft's topological approach is successful, they aim to have a **fault-tolerant prototype (FTP) in "years, not decades"** according to their DARPA-funded timeline [15]– implying possibly late 2020s. A topological quantum computer could, in theory, need far fewer physical qubits per logical qubit, speeding the roadmap to a million physical qubits (because each logical might only need 10-100 physical if topological qubits are naturally protected).
- **Quantum error correction at scale**: Mid-term will refine QEC codes, perhaps moving beyond surface code to more efficient codes if breakthroughs occur (like quantum LDPC codes).
- Techniques like logical qubit networking (connecting logical qubits via teleportation) might be demonstrated, enabling modular scaling beyond a single cryostat or device.
- We will likely see **automation and error mitigation** bridging to error correction. By 2026+, calibrations and operations might be largely automated by AI-driven procedures, which is necessary as system size grows.

By 2030, if goals are met, we could see on the order of **1,000 logical qubits** available with error rates low enough to run medium-depth algorithms (like perhaps factoring a 1024-bit number, or simulating medium molecules with chemical accuracy, etc.). This would require perhaps $\sim$1e5-1e6 physical qubits depending on error rates – essentially the million scale is reached by around 2030 in optimistic scenarios.

Long Term (2030-2035): Full-Scale Fault-Tolerant Quantum Computers and Utility-Scale Integration. In this horizon, the expectation is to **scale from hundreds to many thousands of logical qubits** and to reduce overhead as physical qubits improve. IBM has mentioned a "much larger system by 2033" beyond their 2029 machine [13] – potentially this could mean an order of magnitude more logical qubits or much lower error rates enabling deeper circuits.

Around this time, networks of quantum computers might emerge – linking multiple million-qubit machines to effectively increase capacity or allow distributed quantum computing for even larger problem instances.

Concrete targets:

- PsiQuantum's vision is a million-qubit photonic quantum computer achieving fault tolerance by the mid-2020s to early 2030s. They opted not to build interim small devices, jumping to an "utility-scale" system [16]. The timeline of their Illinois facility suggests within a 5-7 year timeframe (by ∼2030) they aim to have at least a prototype of that large-scale photonic machine operational [16].
- Google might by 2029 have demonstrated error-corrected modules and then work on expanding number of logical qubits through the early 2030s.
- By 2035, one could imagine a fully operational million-qubit commercial quantum computer available via cloud, solving problems that are otherwise impossible. This might be the point where RSA-2048 truly becomes breakable in practice (given the resource estimates of 1 million qubits for 1 week factoring [14]).

To realize this, intermediate technical tasks include:

- Improving qubit quality (coherence and gate fidelity) further so that overhead per logical qubit decreases (maybe we only need 100 physical per logical by then, not 1000).
- Developing **quantum multiplexing**: e.g., one control line controlling multiple qubits sequentially or in frequency space, to reduce control complexity.
- Full integration of cryo-electronics: by the 2030s, the control systems that were external in the 2020s may be fully integrated in cryostat, looking more like a single big machine rather than a physics experiment.
- **Error-aware compilers**: software that optimally schedules logical operations to minimize the space-time volume cost on the QEC lattice. This will allow making the most of limited logical qubits.

Beyond 2035, scaling to tens of millions or a billion qubits might be considered for tackling problems like very high accuracy quantum simulations or global-scale cryptanalysis, but that's beyond our current scope. The key is reaching the first million, which should unlock a broad range of applications.

B. Timelines from Key Players

We can glean specific timelines from public statements:

- **Google**: "useful, error-corrected quantum computer by 2029" [12]. Their roadmap (from Google Quantum AI) had six milestones, culminating in a large-scale error-corrected machine. They have built a new campus for this push and plan to spend billions this decade [12]. By implication, a million-qubit device (since they explicitly mention needing 1,000 physical per logical and targeting 1,000,000 physical) is aimed for 2029 [12]. We might expect Google to have progressive increases: e.g., ∼100 qubits (supremacy) in 2019, ∼1000 qubits by mid-2020s, then multi-chip modules reaching 1e6 by 2029 through aggressive parallelization of chips and a heavy error-correction emphasis.
- **IBM**: IBM's public roadmap has shifted from emphasizing raw chip qubit counts toward modular, fault-tolerant systems. IBM introduced IBM Quantum System Two as a modular platform built around multiple "Heron"-generation processors [26]. In June 2025, IBM stated a goal of a "practical" fault-tolerant quantum computer in 2029 (codenamed **Starling**) and a larger system by 2033 [1]. IBM's 2029 target is often described in terms of *logical* qubits (e.g., ∼200 logical), implying physical-qubit needs that depend strongly on code choice, cycle time, connectivity, and decoder architecture; million-qubit-scale physical systems may plausibly align with the 2030s depending on these assumptions.

Given all these, a reasonable consensus timeline is:

- ∼**2025**: first logical qubit that lives longer than physical (achieved).

- **2025-2027**: chip-level integration of $\sim$1000 qubits, small error-corrected algorithms (e.g. a logical qubit memory with minutes coherence).
- **2028-2030**: prototypes of fault-tolerant quantum computers with $\mathcal{O}(10^5\text{-}10^6)$ physical qubits, demonstrating algorithms impossible for classical computers (e.g., chemistry simulation beyond classical, solving hard optimization with better scaling, etc.). The term "quantum advantage" shifts to "quantum supremacy" in practical tasks.
- **2030-2035**: refinement and scale-out, resulting in full million-qubit machines reliably operating and being deployed for broad use via cloud or dedicated use (government labs, etc.). Possibly multiple competing implementations (superconducting vs. photonic vs. topological) reach this scale, providing diversity in approaches.

C. Human and Capital Resources

Achieving this roadmap is expensive and manpower-intensive. It is often compared to a "space race" or national-scale program. Here are some estimates and considerations:

- **Investment**: Tens of billions of dollars over the next decade globally. The U.S. National Quantum Initiative has provided a framework, but much investment is private (Google, IBM, Microsoft, Amazon, etc., each pouring in billions). PsiQuantum alone raised around $650M; IBM's decade plan will cost billions (they built a $20M+ fridge, new facilities, etc.). The semiconductor industry analogy suggests that to build new fabrication processes (like for spin qubits or photonic chips), costs run in the billions. So one can expect on the order of **$10-20B of investment industry-wide by 2030** to reach these goals.
- **Workforce**: Training quantum engineers is a bottleneck. Universities, backed by NSF and others, are establishing quantum engineering programs. A million-qubit project might require thousands of person-years of labor. For example, writing control software for millions of qubits, or calibrating them, might need large teams or new AI automation tools. The workforce will include not just Ph.D. scientists but also technicians, cryo specialists, software developers, and even facility engineers. A conscious effort to build this workforce is underway (e.g., NSF Quantum Leap Challenge Institutes emphasizing workforce).
- **Multi-organization collaboration**: It is likely no single entity can do it alone. Partnerships (like between companies and national labs, or cross-company alliances for standards) may form. We see early hints: e.g., Google partnering with national labs for materials; IBM partnering with universities worldwide; national labs like NIST and MIT Lincoln Lab working on fabrication for superconducting and ion trap components. The roadmap might involve combining expertise (for instance, one group might develop the cryo-CMOS, another the qubit chip, a third the photonic interconnects, etc., integrated in a final system).

Management of such a project may resemble big science projects like CERN or the Apollo program, requiring coordination and project management talent to keep on schedule.

On the capital front, site selection might consider incentives: state or national governments might subsidize building a facility (like Illinois backing the PsiQuantum park, or StarGate in Amarillo, Texas).

We should also note the **geopolitical dimension**: the U.S., EU, China, etc., all want leadership in quantum computing. This likely means multiple parallel million-qubit efforts. China, for instance, has built a National Laboratory for Quantum Information in Hefei and may aim for similar milestones. Europe's Quantum Flagship is funding companies and consortia toward scalable devices. So by 2030 we may see multiple million-qubit machines globally, not just one.

D. Preferred Geographic Site Characteristics

This overlaps with Section III-C but we can summarize specific characteristics:

- Access to reliable, high-capacity power (possibly near a power plant or substation, or own generator).
- Proximity to academic and industry quantum talent (e.g., near universities or tech hubs like Silicon Valley, Boston, etc., or near national labs).
- Political support and funding (areas where government is investing in high-tech).
- Low natural disaster risk: ideally not in floodplain (unless mitigated), low seismic zone (or building designed for it). For instance, if on West Coast (earthquake-prone), extra cost to earthquake-proof the facility.
- Space for expansion: as technology evolves, more machines or bigger machines may be added, so a campus-like environment is helpful.
- Possibly synergies with semiconductor fabs (for spin qubits or photonic chips, being near a fab like in Arizona, NY, or Belgium (imec) could reduce logistics).
- Climate could be a factor: a cooler climate might slightly ease cooling loads, but with modern HVAC this is secondary. Still, many data centers cluster in temperate climates or places with cheap hydroelectric power (like Pacific Northwest).
- **Security**: quantum computers, especially if used for cryptography tasks, will be strategic assets. Sites might be on secured campuses or have cybersecurity and physical security measures akin to those at supercomputing centers or sensitive research labs.

Examples of sites:

- Poughkeepsie, NY (IBM): has power infrastructure, existing IBM fab nearby, low seismic, and NY state support.
- Santa Barbara, CA (Google): chosen for historical reasons (research group there), but not ideal for scale due to seismic risk; Google might eventually build in a more secure spot for a full data center (maybe at their Iowa or Tennessee data centers, etc., hooking quantum hardware to their cloud).
- Chicago, IL (PsiQuantum): large metro, strong universities, state support, repurposing industrial land – likely robust power from grid and room to build a custom facility.
- Quantum labs at national labs: e.g., Oak Ridge, Argonne (both are in central US with lots of land, good power from TVA or grid, and already host supercomputers, which means infrastructure exists for high power and cooling).

In conclusion, the roadmap to a million qubits is a complex interplay of advancing qubit technology, integrating massive engineering systems, and coordinating efforts across disciplines. The milestones laid out by major players give confidence that if progress continues, the late 2020s to early 2030s will see the first machines hitting the million-qubit scale (in physical qubits), accompanied by enough error correction to perform revolutionary computations.

Achieving this will require not only scientific innovation but also **project management at scale**, significant financial investment, and supportive infrastructure. The preferred sites and setups are already being put in place to make this ambitious goal a reality.

V. Research and Development Landscape

The push toward large-scale quantum computing is supported by a broad ecosystem of private companies, academic institutions, and government programs. Here we summarize the contributions and roles of each:

A. Private Companies and Industry Efforts

A number of high-profile companies are leading quantum computing R&D, each often focusing on a particular technology:

- **IBM**: IBM has one of the most advanced superconducting qubit programs. They operate IBM Quantum, offering cloud access to multiple superconducting processors alongside full-stack tooling (Qiskit). IBM has demonstrated processors such as Eagle (127 qubits), Osprey (433 qubits), and Condor (1121 qubits), and introduced the "Heron" generation optimized for higher fidelity; more recent emphasis is on modular scaling via IBM Quantum System Two and on fault-tolerance milestones (Starling 2029; larger system by 2033) [26], [1]. IBM invests heavily in the full stack: they develop control electronics, the Qiskit software stack, and have research into error correction and applications. They also built the largest cryostats (Project Goldeneye) in anticipation of scaling [17]. IBM's approach is somewhat open – they collaborate with universities (e.g., a large partnership with University of Tokyo and others to form quantum hubs) and publish many results. They emphasize metrics like Quantum Volume and have demonstrated quantum advantage-type experiments (e.g., 127-qubit circuits). IBM also engages with enterprise clients to explore use-cases on smaller machines, preparing a market for the big machine in the future.
- **Google (Alphabet)**: Google's Quantum AI division uses superconducting qubits (and also researches other types like spin). They achieved the 2019 quantum supremacy milestone. Now their focus is on error correction; in 2023 they showed scaling of surface codes. Google is aiming for an error-corrected quantum computer by 2029 [12] and built a state-of-the-art campus with fabrication and labs in Santa Barbara [12]. Google's strategy involves custom hardware (they design their own qubit chips and cryo hardware), but they have not offered public cloud access widely (except some small experiments via Google Cloud). They do collaborate on algorithms (e.g., with Volkswagen on traffic flow optimization experiments). Their efforts are also supported by partnerships (NASA's Ames Research Center was involved in verifying their 2019 result). Google's recent contributions to theory (like resource estimation for RSA) show they are covering both hardware and algorithm readiness [14].
- **Microsoft**: Microsoft took a unique route focusing on topological qubits (Majorana-based). After years of fundamental research, in 2022-2023 they finally had experimental proof of Majorana zero modes and in 2025 announced a prototype chip (Majorana 1) [15]. Microsoft's approach, if it works, could leapfrog in scalability because a topological qubit might be inherently protected, dramatically reducing error rates. Microsoft also developed a strong software ecosystem (Azure Quantum cloud, Q# language, resource estimation tools). They run cloud services that allow use of competitor hardware too (Honeywell/Quantinuum's ion traps, QCI's hardware, etc., via Azure). With DARPA's US2QC program funding, Microsoft is in a race to prove a small fault-tolerant topological qubit array within a few years [15]. They also integrate with their classical cloud (e.g., providing Jupyter notebooks with quantum simulators and linking quantum subroutines with Azure's classical compute). If Microsoft's physics bet pays off, they'll be a major player in delivering the first million-qubit chip.
- **Amazon (AWS)**: Amazon doesn't build its own hardware (yet), but through AWS Braket they provide cloud access to multiple quantum devices (D-Wave annealers, IonQ and Oxford ion traps, Rigetti superconductors). They have a quantum R&D center (AWS Center for Quantum Computing at Caltech) working on superconducting qubits and cryo communication (they published work on connecting two cryostats with microwave links). Amazon likely is preparing to integrate quantum computing as a cloud service seamlessly once hardware is ready. They also invest in error mitigation and compilation (their open-source toolkit PennyLane is popular for variational algorithms). While Amazon's timeline is not public, they are hiring top academics and could surprise with hardware developments in coming years.
- **Intel**: Intel is focusing on silicon spin qubits, leveraging their semiconductor manufacturing prowess. They have fabricated qubit chips on 300 mm wafers and developed a cryogenic control chip "Horse Ridge"

to operate spin qubits at low temperature. Intel collaborates with QuTech (Netherlands) for research [3]. Their advantage is making devices using advanced transistor fab techniques, which could yield good uniformity and eventually high-density qubit arrays. In 2022, Intel demonstrated 2-qubit gates in silicon and is steadily improving fidelity. Intel's long-term vision is to apply the techniques of Moore's Law (CMOS scaling) to quantum – meaning if one spin qubit works, they can try to scale to thousands by leveraging chip design know-how. Intel's timeline has been less aggressive in public, but they quietly aim for large-scale spin-qubit processors in the 2030s. They recently stated error rates they achieved are at the threshold for fault-tolerance, indicating they are overcoming major hurdles.

- **Quantinuum (Honeywell + Cambridge Quantum)**: Quantinuum operates some of the highest-performing trapped-ion systems (the System Model H1 with 20 qubits, achieving record quantum volume 2^{20}). They focus on high fidelity and quantum software (Cambridge Quantum brings algorithm expertise, e.g., in chemistry with their TKET compiler). Quantinuum likely will scale via ion trap architectures with shuttling or photonic links. They haven't given a specific million-qubit date, but they are considered a front-runner in near-term quantum computing due to very low error rates (one of their devices achieved an error rate 1e-4, and they have demonstrated repetitive error correction on a logical qubit). Their strategy is to sell solutions (encryption keys with quantum random, etc.) now and build toward larger systems steadily.
- **IonQ**: A spin-off from University of Maryland and Duke, IonQ has small ion traps (11-32 qubits) and went public via SPAC in 2021 to raise funding. They are expanding capacity and have roadmapped a 64-qubit device by 2025 and modular scaling beyond. IonQ's approach includes new trap designs and perhaps integrating photonic coupling of multiple traps. They are also exploring barium ions for better stability. With significant capital, IonQ is one of the few pure-play quantum companies focusing solely on building a universal quantum computer with high qubit count. They might aim for a few hundred qubits by 2027 and use error mitigation or minimal error correction to do useful tasks, then rely on networking for larger scale.
- **Rigetti Computing**: Rigetti is a smaller player focusing on superconducting qubits. They built 30+ qubit chips and aim for modular 80-qubit tiles. Rigetti's distinguishing factor is pursuing multi-chip scaling (they demonstrated a prototype with two chips connected by aluminum wirebonds entangling qubits across chips). They also develop their own control hardware and Quil programming language. Rigetti has had technical and financial challenges recently, but they still contribute to innovation in packaging and cloud deployment of superconducting qubits.
- **D-Wave Systems**: D-Wave specializes in quantum annealing, not gate-model computing (though they have announced a small gate-model effort too). D-Wave's annealers have $>$5000 qubits (albeit flux qubits that are not individually gate-addressable but used for annealing tasks). They have a roadmap to $\sim$7000 qubits on their next annealer (Advantage2). While annealing is a narrower paradigm, D-Wave's advancements in cryogenic integration (thousands of qubits with interconnects on a chip) and fabricating that many superconducting devices are notable. If someday annealing and gate-model converge or annealing finds niche uses, D-Wave could have the first "thousands of qubits" quantum machines doing something useful, albeit not general algorithms. They also now provide cloud access and have a user base exploring optimization problems.
- **Xanadu**: A startup focusing on photonic quantum computing using continuous variables (squeezed light). Xanadu built a photonic chip called Borealis that did Gaussian Boson Sampling with 216 squeezed-state qubits, achieving a quantum advantage claim in 2022. They also develop PennyLane software for quantum ML. Xanadu's goal is to create a photonic gate-based machine; they are working on squeezing-based qubits (gkp qubits potentially). They have significant funding and a roadmap to a fault-tolerant photonic architecture (likely also cluster-state MBQC). They might be a competitor to PsiQuantum in photonics.
- **PsiQuantum**: Already discussed, PsiQuantum is all-in on photonic fusion-based quantum computing.

They've partnered with GlobalFoundries to fabricate photonic chips in standard CMOS processes [16]. They are secretive about progress, but their collaboration to build a facility hints that they are now moving from R&D to implementation phase. PsiQuantum has stated that small NISQ machines won't deliver value; they aim directly for million-qubit fault-tolerant device. Their team's experience (many from optical fiber telecom background) and large funding make them a serious contender for achieving a million-qubit system in the early 2030s.

- **Others**: There are many startups globally: **Silicon Quantum Computing (Australia)** working on donor spins and atomic-scale qubits, **Quantum Circuits Inc. (QCI)** focusing on 3D superconducting cavities, **Alice & Bob (France)** developing superconducting cat-qubits (bosonic qubits with intrinsic error bias), **Pasqal (France)** and **Atom Computing (US)** with neutral atoms, **ORCA Computing (UK)** on photonic quantum memory and small processors, **Quantum Brilliance (Australia/Germany)** on room-temperature diamond NV systems, etc. Each of these contributes innovations. For example, Alice & Bob's "cat qubit" showed that passively stabilized qubits might reduce error correction overhead by autonomously correcting bit-flips. If such approaches succeed, they could drastically cut the number of qubits needed for fault tolerance, indirectly aiding the million-qubit goal.

In summary, the private sector is vibrant. **Competition and collaboration** coexist: IBM, Google, Intel, etc., often collaborate with academia and government on basic research, even as they compete for milestones. This competitive drive helps push the field forward rapidly.

B. Academic Efforts

Academic research groups and institutions continue to be at the heart of many breakthroughs:

- Universities pioneered many platforms: Ion traps (NIST/University of Colorado, Innsbruck, Oxford), superconducting qubits (Yale, UC Santa Barbara, ETH Zurich), neutral atoms (Harvard, Institut d'Optique, Caltech), silicon spins (Delft TU, UNSW Sydney, Princeton). These groups still produce leading results (e.g., the first 0.9999 fidelity two-qubit gate in ion traps was at NIST; the first 0.9999 fidelity spin qubit gates at Delft; the surface code experiments at Google had academic collaborators).
- Many universities have **Quantum Centers** now: e.g., MIT has the Center for Quantum Engineering, Harvard-MIT Center for Ultracold Atoms (neutral atom QC), University of Maryland hosts the Joint Quantum Institute and Joint Center for Quantum Information and Computer Science (in partnership with NIST), which have done a lot on ion traps and modular systems. University of Maryland in particular gave birth to IonQ and continues to do advanced trap research (recently demonstrated programmable 2D ion arrays, etc.). Duke University is another ion trap hub (its research led to the IonQ architecture).
- In Europe, **University of Oxford** and **University of Innsbruck** are ion trap powerhouses (Innsbruck's team under Rainer Blatt and Thomas Monz has set records in multi-ion entanglement, Oxford's group under David Lucas and others has world-leading fidelity). **TU Delft** and **U. of Copenhagen** lead in spin qubits and Majorana (in collaboration with Microsoft for Majorana; Delft QuTech with Intel for spins). **ETH Zurich** and **Karlsruhe Institute of Technology** do superconducting circuits research.
- Academic consortia have formed under national programs: for example, the **U.S. Quantum Leap Challenge Institutes (QLCI)** funded by NSF include institutes like Q-SEnSE (quantum sensors but related tech), HQAN (Hybrid Quantum Architectures and Networks) which links cold atoms with superconducting systems (Univ. of Illinois/UIUC, Wisconsin, etc.), and **Software-Tailored Architecture for Quantum co-design (STAQ)** that involved multiple schools to build a small ion-trap quantum computer.
- The **US Department of Energy Quantum Centers** (funded since 2020) involve universities and labs: e.g., **SQMS (Superconducting Quantum Materials and Systems Center)** led by Fermilab, focusing on high-quality superconducting resonators and qubits (with partners like Northwestern, Ames Lab, etc.);

QSA (Quantum Systems Accelerator) led by Lawrence Berkeley and Sandia, focusing on ion and atom systems (with partners like Berkeley, MIT, Univ. Maryland, Duke etc.); **Q-NEXT** led by Argonne focusing on quantum interconnects and materials (including University of Chicago and others).

- These academic efforts contribute pieces of the puzzle: materials science (making better qubits and reducing noise), new gate techniques, novel error-correcting codes (academics propose many new codes and error mitigation strategies), and algorithms that could reduce hardware demands.

One should also acknowledge the importance of **National Metrology Institutes** like NIST (US) and PTB (Germany) in quantum computing research – NIST provided foundational work in ion traps and superconducting circuit metrology. They continue to refine techniques for high-fidelity control.

In academia, there is also a lot of **cross-pollination**: e.g., a theory group might work with a company to test an algorithm on actual hardware (like a university algorithm team running on IBM machines via the cloud). And academic labs often try more radical ideas that companies might not prioritize due to risk – e.g., novel qubit types (photonics with orbital angular momentum, or Rydberg arrays in new geometries).

Academic research provides the training ground for the workforce. Many startup founders and industry researchers came out of PhD programs that specialized in QC. The academic freedom allows exploration of next-gen ideas that could become crucial (like the cat qubits from Yale that led to startups like Alice&Bob).

C. U.S. Government Programs

The U.S. government has a multi-pronged support system for quantum computing R&D:

- **National Quantum Initiative (NQI)**: Enacted in 2018, it coordinates funding across NSF, DOE, and NIST (and other agencies like DoD). It established NSF QLCI centers and DOE QIS centers (as described above) and expanded NIST work. It also set up the National Quantum Coordination Office (NQCO) and an advisory committee. The NQI provides a strategic vision and has authorized ∼$1.2 billion over 5 years (2019-2023) across agencies.
- **NSF Programs**: Aside from QLCIs, NSF has smaller programs like Quantum Foundries for materials, and core programs in quantum information science. The NSF also funds workforce development, e.g., traineeships in Quantum Engineering.
- **Department of Energy (DOE)**: DOE's Office of Science funneled significant funding into the five National QIS Research Centers (like SQMS, QSA, etc.), each funded at ∼$25M/year for 5 years [18]. These centers involve national labs, universities, and sometimes industry. DOE labs also individually pursue quantum computing in their mission areas (e.g., Oak Ridge and Lawrence Berkeley run quantum computing user programs using IBM and other hardware [18]; Los Alamos does a lot of quantum algorithms research; Sandia develops ion trap hardware; Argonne works on superconducting and spin qubit materials). DOE's emphasis is often on how quantum computing can help science (simulation of materials, chemistry, nuclear physics) and also on the underlying technology (fabrication, new materials like topological superconductors, etc.).
- **Department of Defense (DoD)**: DoD agencies invest in quantum computing both for potential future computing needs and to avoid being surprised technologically.
  - **DARPA (Defense Advanced Research Projects Agency)**: DARPA launched the ONISQ (Optimization with NISQ) program around 2020 to explore near-term uses of quantum for optimization (with participants like Google, NASA, and some startups). DARPA's Quantum Benchmarking program tries to quantify quantum advantage rigorously. Most notably, DARPA's US2QC (Underexplored Systems for Utility-Scale Quantum Computing) program started in 2022 funded companies (report-

edly Microsoft, Atom Computing, and PsiQuantum) to pursue out-of-the-mainstream approaches to million-qubit systems [18]. Microsoft's topological qubit and PsiQuantum's photonics are likely under this program's umbrella (Atom Computing is working on nuclear spins or neutral atoms as a dark horse). DARPA is giving these companies a few years and several phases to prove viability and eventually deliver a prototype. This injects significant funding (exact amounts classified, but DARPA programs of this size might be tens of millions each) to non-superconducting approaches specifically. DARPA is also funding enabling tech like the Atomic-Photonic Integration (A-PhI) program which developed on-chip optical trapping of atoms [18], relevant for clocks and potentially neutral atom QC, and other quantum sensing programs that can spill over to computing tech (e.g., better lasers, better quantum memories).

  - **IARPA (Intelligence Advanced Research Projects Activity)**: IARPA (under ODNI) has been investing in quantum computing for a while. The **LogiQ** program (2016-2021) aimed to demonstrate a logical qubit (focused on ion traps) [24]. It funded groups like Chris Monroe's at UMD and Honeywell, and reportedly they did achieve a simple logical qubit with some error correction [25]. A successor, **ELQ (Entangled Logical Qubits)**, started around 2021 to take it further (maybe multiple logical qubits entangled). IARPA also had programs on specialized quantum computing: e.g., QEO (quantum enhanced optimization) looked at quantum annealing and analog approaches; and there are connections to quantum algorithms for chemistry (through IARPA's molecular simulation initiatives). The intelligence community is interested because of cryptography – they track progress to anticipate when a quantum computer could break encryption. The recent progress (1 million qubits for RSA in a week) [14] underscores this is not imminent tomorrow but also not centuries away, hence sustained funding.
  - **Laboratory Research**: The Army Research Office (ARO) and Air Force Office of Scientific Research (AFOSR) have long funded basic research in quantum computing in academia. ARO's Ion Trap program helped develop some leading ion technologies; AFOSR funded superconducting qubit physics. Army Research Lab and Naval Research Lab have internal efforts too (NRL did work on photon sources, ARL on benchmarking). The Department of Defense also established e.g. an Army-led Quantum Research Center with academic partnerships to look at long-term quantum tech.

- **National Institute of Standards and Technology (NIST)**: NIST is a smaller but highly influential player. NIST physicists (like Dr. John Martinis before he went to Google, or current groups) contribute to fundamental metrology of qubits, developing high-accuracy measurements, new techniques like ion trap chip technology, etc. NIST co-runs the JQI and JILA (with Univ. of Colorado) which pioneer many quantum computing building blocks. NIST also helps set standards – for example, they worked on a common quantum assembly language (OpenQASM with IBM). NIST's recent main contribution to the community is leading the post-quantum cryptography standardization (to secure communications against quantum attacks in advance) [14], which while not building the quantum computer, is critical to mitigate its eventual capabilities in cryptography.
- **Other Agencies**:
  - **The National Security Agency (NSA)** historically funded a lot of early quantum computing (particularly superconducting) research under programs like "Penetrating Hard Targets," though details are classified.
  - **The Department of Energy's NNSA (National Nuclear Security Administration)** might have interest in quantum computing for simulating physics and also breaking encryption – they fund some of the DOE lab work.
  - The **National Science Foundation (NSF)**, as mentioned, funds academic research broadly. It's worth noting an NSF center, **IQIM at Caltech (Institute for Quantum Information and Matter)** which is a long-standing NSF Physics Frontiers Center, has produced top-notch theoretical work (many

quantum algorithms and error correction schemes come from there or similar centers like KITP at UCSB).
- **NASA**: NASA's interest is in using quantum computing for optimization and simulation relevant to aeronautics and space. NASA Ames collaborated with Google on the supremacy experiment (providing supercomputing resources to verify classical simulation bounds). NASA also had a Quantum Artificial Intelligence Lab (QuAIL) that looked into D-Wave annealers for scheduling problems. NASA remains a potential end-user for large quantum computers for tasks like mission planning optimization or even quantum chemistry (for fuel and materials).
- **DOE Office of Naval Research (ONR)**: funds some quantum research (like topological qubits, quantum sensing).
- The **White House OSTP** has been supportive, releasing quantum strategy documents, and the US Congress has allocated increasing budgets.

Collectively, U.S. government support ensures long-range research that may be too risky or foundational for private companies alone. It also encourages collaboration to avoid duplication of effort and to push the entire field toward common goals (like fault-tolerance).

Globally, it's similar: Europe has flagship programs, and individual countries (UK, Germany, France, Netherlands, Australia, China, etc.) have their own government-funded initiatives. For example, China reportedly invested heavily in superconducting qubits (with groups achieving 66-qubit devices and photonic advantage experiments under USTC), and launched a Quantum Initiative of their own. Europe's Quantum Flagship (billion-euro program) funds many consortia (like OpenSuperQ for superconducting, AQTION for ion traps, etc.). These all contribute to the knowledge base and sometimes work in partnership with U.S. efforts (quantum science is still quite collaborative internationally, though there is some competition and concerns about technology leadership).

Finally, government agencies also help with **standards and consortia**: for instance, NIST and the EU have organized workshops on quantum benchmarks, interoperability (like quantum internet protocols), and even foundries (the U.S. has the Quantum Foundry at UCSB for superconducting circuits, while EU has efforts to coordinate fabrication through projects like QuoVidQC, etc.).

Summary of Support: The drive to million-qubit quantum computers is bolstered by:

- **Private sector innovation and funding**, bringing engineering discipline and large-scale integration.
- **Academic research and education**, providing breakthroughs in theory and prototypes as well as training the workforce.
- **Government programs and funding**, ensuring long-term and high-risk research is pursued, coordinating national strategies, and protecting security interests.

This triple-helix of industry, academia, and government is what has brought quantum computing from a scientific curiosity in the 1990s to the brink of technological reality in the 2020s. Each plays a role:

- Universities invent new methods and educate new scientists.
- Companies build devices and push towards commercialization.
- Governments fund infrastructure and address areas that are beyond the scope of individual companies (like fundamental theory, or bridging the "valley of death" from lab demo to technology).

As we stand today, all these efforts together have led to the first steps of quantum error correction, the expansion of qubit counts, and concrete plans for scaling. With continued support, the ambitious goal of a million-qubit quantum computer – once seen as purely theoretical – is now believed to be achievable within

a decade or so [12], [13], ushering in the era of quantum advantage for real-world problems.

Tables II and III below provide an illustrative summary of some major organizations and their quantum computing focus.

VI. Conclusion

In conclusion, the race to build a million-qubit quantum computer is a collaborative and competitive effort spanning private tech giants, cutting-edge academic labs, and strategic government programs. The **state of the art** today showcases multiple qubit technologies each with strengths: superconducting circuits and trapped ions lead in maturity and fidelity, while neutral atoms, photonics, and spins promise greater scalability. **Theoretical advances** in error correction and algorithms are guiding how these qubits can be scaled to useful quantum computing.

The **physics** underpinning each platform – from Josephson junctions enabling superconducting qubits [20], to electromagnetic trapping of atomic ions [8], optical tweezers arranging neutral atoms [9], and photons carrying quantum states through optical networks [10] – is well understood, giving confidence that they can be engineered into larger systems.

From a **systems engineering** perspective, we have identified the formidable challenges of power, cryogenics, and control for scaling to a million qubits, and described emerging solutions like cryogenic integration and modular architectures [3]. Building such a machine demands facility-scale planning, potentially borrowing solutions from classical supercomputing infrastructure (and even considering on-site power generation for reliability).

A plausible **roadmap** has been outlined, projecting that by the end of this decade we will see the first fault-tolerant quantum computers with on the order of 10^5-10^6 physical qubits [12], [13], and by the early-to-mid 2030s, fully realized million-qubit systems capable of transformational computational feats. This will require sustained investment and coordination – akin to a "moonshot" – including training hundreds of experts and spending billions of dollars across public and private sectors.

Finally, we reviewed the rich **R&D ecosystem** supporting this effort. Industry leaders like IBM, Google, and others drive technology integration and set aggressive milestones [3], [12]. Academic institutions supply new ideas and talent, often supported by government grants. U.S. government programs (mirrored by efforts in other countries) ensure that fundamental roadblocks are addressed and that the nation remains at the forefront of quantum technology development – not least because of the profound economic and security implications of success in this field.

In summary, the quest for a million-qubit quantum computer is well underway. Each qubit modality – whether a tiny Josephson junction circuit or an atom hovering in laser light – adds to the growing quantum toolkit. Advances in error correction are converting fragile qubits into reliable logical units [14]. Engineering innovations are solving the practical challenges of control and scale [3]. If progress continues apace, the coming decade will likely witness the emergence of quantum computers that can tackle problems utterly beyond the reach of classical supercomputers, fulfilling the long-held promise of quantum computation to revolutionize fields from cryptography to chemistry and beyond. The convergence of **scientific insight and engineering prowess**, supported by a broad coalition of researchers, companies, and governments, is making what once seemed a fantastical goal – a universal quantum computer – into a foreseeable reality.

TABLE II

SELECT MAJOR QUANTUM COMPUTING R&D PLAYERS AND THEIR FOCUS.

Organization	**Type**	**Quantum Technology Focus**	**Notable Milestones / Programs**
IBM	Private Company	Superconducting qubits, full-stack software (Qiskit)	Eagle/Osprey/Condor-era devices; modular System Two platform (Heron generation); Starling 2029 and larger system by 2033 (fault-tolerant roadmap)
Google Quantum AI	Private Company	Superconducting qubits (Sycamore), quantum algorithms	Achieved quantum supremacy (53 qubits, 2019); Demonstrated improving QEC (2023); Aiming for useful error-corrected computer by 2029
Microsoft	Private Company	Topological qubits (Majorana), Azure Quantum software	"Majorana 1" program announced (2025) with long-term scaling roadmap; high-variance pathway pending repeatable devices and error-rate characterization
Intel	Private Company	Silicon spin qubits, cryo-CMOS control	2018-2020: 49-qubit superconducting test chip; Focus shifted to spin qubits: demonstrated 99.9% fidelity in Si double-dot (2022); 300 mm fab process for qubits
Quantinuum	Private (Honeywell + Cambridge)	Trapped-ion qubits, quantum software	High-fidelity ion systems (50+ physical qubits reported); strong benchmarking metrics (e.g., Quantum Volume) and early logical-qubit demonstrations
IonQ	Private Company	Trapped-ion qubits (Yb, Ba); modular traps	11-qubit system (2019); 32-qubit system (2021); planning networked traps (64+ qubits by 2025)
Rigetti	Private Company	Superconducting qubits, multi-chip modules	8-qubit (2017), 31-qubit (2019) processors; developing 80-qubit modular chip; went public 2022
D-Wave	Private Company	Quantum annealing (flux qubits); some gate-model R&D	5000+ qubit annealer Advantage (2020); exploring 2-qubit gate on flux qubits (small-scale project)
PsiQuantum	Private Company	Photonic qubits (silicon photonics, cluster-state)	No public device yet; Partnered with GlobalFoundries; Plans for utility-scale (1M qubit) machine in Illinois

ACKNOWLEDGMENTS

The author thanks colleagues across industry and academia for discussions and publications that informed this work.

TABLE III

SELECT MAJOR QUANTUM COMPUTING R&D PLAYERS AND THEIR FOCUS.

Organization	**Type**	**Quantum Technology Focus**	**Notable Milestones / Programs**
Xanadu	Private Company	Photonic qubits (continuous-variable, squeezing)	Demonstrated 216-mode Gaussian Boson Sampling (2022) with quantum advantage claim; Developing fault-tolerant photonic architecture
University of Maryland / Duke (Ion Trap groups)	Academic / Startup (IonQ roots)	Trapped-ion gate techniques, photonic linking	Developed key ion gate protocols (Molmer-Sorensen); first ion-qubit networking demos
Yale University	Academic (Many startups spun out)	Superconducting qubits (3D cavities, cat qubits)	Invented transmon qubit ($\sim$2007), bosonic error-correcting codes ($\sim$2016); basis of startups like Quantum Circuits Inc. and Alice&Bob (cat qubits)
Delft University of Technology (QuTech)	Academic / Govt partnership	Spin qubits, topological qubits (Majorana)	Demonstrated 2-qubit logic in silicon MOS dots; key partner with Intel and Microsoft; first detection of Majorana modes (2018, revised 2022)
NIST (National Inst. Standards & Tech)	Government (research lab)	Ion traps, superconducting circuits, metrology	NIST ion traps achieved record fidelities; co-developed surface code techniques; NIST teams won Nobel prizes for ion control. Leading post-quantum crypto standards
DOE National Labs (e.g. Fermilab, Lawrence Berkeley, Sandia)	Government (DOE Science)	Various (SC qubits at Fermilab, neutral atoms & ions at LBL/Sandia, networking at Argonne)	Established QIS Centers (SQMS – high-Q cavities; QSA – trapped ions/atoms; Q-NEXT – interconnects, etc.). Sandia developed high-optical-access ion traps; LBL advancing error correction codes
DARPA (US DoD)	Government (Defense R&D)	Program funding for novel QC approaches	ONISQ (2020-) for NISQ optimization; US2QC (2022-) funding 3 companies (e.g. PsiQuantum, Microsoft) for million-qubit concepts; various quantum tech programs (e.g. A-PhI, QuDAC)
IARPA (US Intel)	Government (Intel R&D)	Funding for error-correction, alt platforms	LogiQ program (2016-21) achieved logical qubit demo; Follow-on ELQ program; also investing in quantum algorithms.

NOTE ON REFERENCES

While there are many research papers that span these topics, this article emphasized papers with the following metrics:

- ✓ High Citation Counts – Many references have 1,000+ citations, ensuring they are widely accepted and influential in the Quantum, Math, Policy, Economics, and National Security research community.
- ✓ Top Quantum, Math, Policy, Economics, and National Security Institutions – Includes research from Stanford, CSET, GWU, and Harvard.
- ✓ Authors with High h-Index – leading AI researchers such as Craig Gidney (h-index: 150+).

REFERENCES

[1] "IBM aims for quantum computer in 2029, lays out road map for larger systems," Reuters. https://www.reuters.com/business/retail-consumer/ibm-aims-quantum-computer-2029-lays-out-road-map-larger-systems-2025-06-10/

[2] "Trump signs legislation to boost quantum computing research with $1.2 billion," GeekWire https://www.geekwire.com/2018/trump-signs-legislation-back-quantum-computing-research-1-2-billion/

[3] Quantum Computing: Whose Qubit is Better? @ ISSCC 2025 https://sscs.ieee.org/event/quantum-computing-whose-qubit-is-better/

[4] Hanhee Paik (IBM) - IBM's 1000-qubit Condor and QEC plans, coherence and fidelity metrics

[5] Shirin Montazeri (Google) - Scaling control electronics to millions of qubits, challenge and opportunity

[6] Sophie Bene (imec) - Silicon spin qubits: long coherence and CMOS scalability

[7] Jeanette Roberts (Intel) - Spin qubits for fault-tolerance, millions of qubits with fab tech

[8] Eleventh Hour Enthusiast (Medium) - Trapped-ion hardware and performance, overview of ion trap advantages (identical qubits, high fidelity, all-to-all connectivity) https://medium.com/@EleventhHourEnthusiast/ion-based-quantum-computing-hardware-performance-and-end-user-perspective-d1b7ab24de7c

[9] Eleventh Hour Enthusiast (Medium) - Neutral atom hardware and performance, Rydberg gates, coherence $\sim$ seconds, fidelities $\sim$99% https://medium.com/@EleventhHourEnthusiast/neutral-atom-quantum-computing-hardware-performance-and-end-user-perspective-98c13456ad23

[10] Quandela Glossary - Photonic qubits, encodings and challenges (single-photon sources, detectors, lack of native two-qubit interactions) https://www.quandela.com/resources/quantum-computing-glossary/photonic-qubit/

[11] Silicon spin qubit advances, 0.5 s coherence, 99.95% 1-qubit and 99% 2-qubit fidelity at 1 K https://phys.org/news/2025-05-silicon-qubits-gain-ground-candidate.html

[12] The Verge (Google AI) - Google's quantum roadmap, targeting 1,000,000 physical qubits by 2029 for error-corrected machine

[13] IBM Quantum Blog - Roadmap to large-scale fault-tolerant quantum computing (through 2033). https://www.ibm.com/quantum/blog/quantum-roadmap-2033

[14] C. Gidney, "How to factor 2048 bit RSA integers with less than a million noisy qubits," arXiv:2505.15917 (2025). https://arxiv.org/abs/2505.15917

[15] Microsoft Azure Quantum Blog (2025) - "Majorana 1" topological-processor announcement and roadmap (high-variance pathway). https://azure.microsoft.com/en-us/blog/quantum/2025/02/19/microsoft-unveils-majorana-1-the-worlds-first-quantum-processor-powered-by-topological-qubits/

[16] "PsiQuantum breaks ground on Chicago quantum site after $1 billion funding," Reuters (Sep. 30, 2025). https://www.reuters.com/business/psiquantum-breaks-ground-chicago-quantum-site-after-1-billion-funding-2025-09-30/

[17] IBM Blog (2022) - Project Goldeneye super-fridge, 1.7 m^3 dilution refrigerator for large experiments, groundwork for scaling

[18] National Quantum Initiative Annual Report FY2024 - DARPA US2QC and other agency efforts, three companies contracted in 2023 to design utility-scale QC

[19] IEEE Spectrum / HPCwire - IBM's 127-qubit and 433-qubit achievements, quantum volume 128->256, plan for 1121-qubit Condor in 2023 https://www.ibm.com/quantum/blog/quantum-metric-layer-fidelity

[20] Spin Quanta - How Josephson Junctions Power Quantum Computing https://www.spinquanta.com/news-detail/the-role-of-josephson-junctions-in-quantum-computing

[21] PsiQuantum "A manufacturable platform for photonic quantum computing" https://www.nature.com/articles/s41586-025-08820-7

[22] "Linear Optics to Scalable Photonic Quantum Computing" https://arxiv.org/html/2501.02513v1

[23] "Alice & Bob to Build $50 Million Advanced Quantum Lab in Paris" https://alice-bob.com/newsroom/alice-bob-product-development-lab

[24] LogiQ (Logical Qubits) https://www.iarpa.gov/research-programs/logiq

[25] "IARPA Pursuing Significant Advancement in Quantum Computing" https://www.odni.gov/index.php/newsroom/press-releases/press-releases-2024/3772-iarpa-pursuing-significant-advancement-in-quantum-computing

[26] "IBM debuts next-generation quantum processor and IBM Quantum System Two, extends roadmap to advance era of quantum utility," IBM Newsroom (Dec. 4, 2023). https://newsroom.ibm.com/2023-12-04-IBM-Debuts-Next-Generation-Quantum-Processor-IBM-Quantum-System-Two,-Extends-Roadmap-to-Advance-Era-of-Quantum-Utility

[27] S. Bravyi *et al.*, "High-threshold and low-overhead fault-tolerant quantum memory," Nature (2024). https://www.nature.com/articles/s41586-024-07107-7

[28] L. Cherry, "Relay-BP: Decoder for qLDPC error correction," IBM Quantum Blog (Aug. 4, 2025). https://www.ibm.com/quantum/blog/relay-bp-error-correction-decoder

[29] P. Webster *et al.*, "The Pinnacle Architecture ..." arXiv:2602.11457 (2026). https://arxiv.org/abs/2602.11457