Quantum Computing in 2026: A Comprehensive Architectural, Physical, and Algorithmic Analysis of the Transition from Noisy Intermediate-Scale Quantum (NISQ) Processors to Fault-Tolerant Logical Qubits, High-Rate qLDPC Codes, and Quantum-Centric Supercomputing

Sandeep, Gautam, Sankalp, Dr. Vishal Shrivastava, Dr.Akhil Pandey, Er.Ram Babu Buri

Department of Computer Science and Engineering, Arya College of Engineering & I.T.
(Affiliated to Rajasthan Technical University, Kota), Jaipur, Rajasthan, India

Received:Oct 06, 2026 Accepted:Nov 03, 2026 Published:Nov 06, 2026

Abstract*—The years 2025 and 2026 mark a historic inflection point in quantum information science: the definitive transition from the Noisy Intermediate-Scale Quantum (NISQ) era toward early Fault-Tolerant Quantum Computing (FTQC) and Quantum-Centric Supercomputing. For nearly three decades, the central impediment to scalable quantum computation was that adding physical qubits increased aggregate decoherence faster than Quantum Error Correction (QEC) codes could suppress it. This barrier was decisively shattered by recent experimental milestones demonstrating below-threshold exponential error suppression on superconducting surface codes (Google's 105-qubit Willow and USTC's 107-qubit Zuchongzhi 3.2), high-rate Bivariate Bicycle quantum Low-Density Parity-Check (qLDPC) memory modules reducing physical qubit overhead by over 90% (IBM's Heron, Nighthawk, and 2026 Kookaburra architectures), reconfigurable neutral-atom optical tweezer arrays executing transversal logical gates (QuEra and Harvard), 99.92%–99.99% fidelity trapped-ion systems (Quantinuum Helios and Oxford Ionics), bosonic cat-qubit processors (AWS Ocelot), and topological InAs-Al nanowire interferometry (Microsoft Majorana 1). This paper presents a comprehensive architectural, mathematical, and algorithmic investigation of the 2026 quantum computing landscape. We formalize a five-layer Quantum-Centric Supercomputing reference stack integrating heterogeneous physical modalities, 4K cryo-CMOS and NVQLink GPU-QPU interconnects, sub-microsecond real-time syndrome decoders, CUDA-Q/Qiskit 2.0 hybrid compilers, and domain applications. We formulate rigorous mathematical models for Lindblad open-system decoherence, surface-code below-threshold scaling (Lambda > 1), Bivariate Bicycle [[144, 12, 12]] qLDPC codes, bosonic cat-qubit bit-flip suppression, and hybrid Variational/Fault-Tolerant Hamiltonian simulation. Finally, we evaluate NIST Post-Quantum Cryptography (FIPS 203 ML-KEM and FIPS 204 ML-DSA) migration imperatives and chart the roadmap toward mega-quop fault-tolerant machines.

***Index Terms*—Bosonic Cat Qubits, Bivariate Bicycle Codes, Fault-Tolerant Quantum Computing (FTQC), Logical Qubits, Neutral Atom Arrays, Post-Quantum Cryptography (PQC), Quantum Error Correction (QEC), Quantum Low-Density Parity-Check (qLDPC) Codes, Quantum-Centric Supercomputing, Superconducting Qubits, Surface Code.**

I. INTRODUCTION

A. The 2025–2026 Inflection Point: From NISQ to Logical Qubits

Since John Preskill coined the term Noisy Intermediate-Scale Quantum (NISQ) in 2018 [1], the quantum computing industry operated in a regime characterized by 50 to 1,000 unerror-corrected physical qubits whose two-qubit gate infidelities (0.1% to 1.5%) and limited coherence times (T_1, T_2 ~ 50–300 microseconds on superconducting transmons) restricted circuit depths to fewer than a few hundred gates before quantum information was thermalized by environmental noise [1], [2]. While NISQ processors achieved landmark demonstrations of quantum supremacy via Random Circuit Sampling (RCS) [3], [4] and utility-scale error-mitigated expectation value estimation [5], unmitigated physical noise prevented the execution of deep, fault-tolerant algorithms such as Shor's prime factorization or Quantum Phase Estimation (QPE) for industrial catalysis [2], [6].

Between late 2024 and 2026, however, a series of foundational hardware and error-correction breakthroughs fundamentally shifted the field from **physical qubit count competition** to **error-corrected logical qubit scaling** [4], [7], [8], [9]:

- **1) Below-Threshold Surface Code Error Suppression:** In December 2024 and throughout 2025–2026, Google Quantum AI demonstrated on its 105-qubit **Willow** superconducting processor that scaling a rotated surface-code logical qubit from a 3x3 grid (distance d=3) to 5x5 (d=5) and 7x7 (d=7) reduced the logical error rate per cycle exponentially by a factor of Lambda = 2.14 per step—achieving a logical lifetime (T_1,L = 291 microseconds) more than double that of the best constituent physical qubit [4]. Shortly thereafter, the Zuchongzhi 3.2 (107-qubit) processor demonstrated

below-threshold error correction using an all-microwave control architecture [10].

- **2) High-Rate qLDPC Codes (90% Overhead Reduction):** Recognizing that planar surface codes require thousands of physical qubits per logical qubit, Bravyi et al. at IBM introduced **Bivariate Bicycle (BB) quantum Low-Density Parity-Check (qLDPC)** codes—specifically the [[144, 12, 12]] 'Gross code'—which encode 12 logical qubits into just 144 physical qubits using a quasi-planar bilayer layout, slashing physical hardware overhead by over 90% compared to surface codes [7]. This architecture forms the foundation of IBM's 2025 **Nighthawk** (120 qubits, 218 tunable couplers) and 2026 **Kookaburra** modular qLDPC logical processors [11].
- **3) Reconfigurable Neutral Atoms and High-Fidelity Trapped Ions:** Bluvstein et al. (Harvard/QuEra) demonstrated fault-tolerant logical algorithms on 48 logical qubits using optically shuttled neutral Rubidium-87 atoms in acoustic-optic deflector (AOD) tweezers [8], while Quantinuum launched **Helios** (November 2025)—a 98-qubit trapped-ion system achieving 99.921% two-qubit gate fidelity, 48 error-corrected logical qubits at a 2:1 encoding ratio, and real-time decoding via NVIDIA GB200 GPUs over NVQLink [9], [12]. Oxford Ionics separately achieved a record 99.99% two-qubit gate fidelity using electronic qubit control [13].
- **4) Hardware-Intrinsic Bosonic Cat and Topological Qubits:** In February 2025, AWS unveiled **Ocelot**, a bosonic processor utilizing Kerr/dissipative 'cat qubits' in superconducting microwave cavities that intrinsically suppress bit-flip errors to ~1 second (>1,000x longer than transmons), reducing QEC overhead by up to 90% [14]. Concurrently, Microsoft unveiled **Majorana 1** (and follow-up **Majorana 2** in 2026), an Indium Arsenide–Aluminum (InAs-Al) 'topoconductor' nanowire array employing radio-frequency (RF) interferometric parity readout toward topological tetron qubits [15], [16].

B. Research Contributions and Paper Organization

This paper provides a rigorous, unified engineering and mathematical treatment of the 2026 quantum computing ecosystem. Our primary contributions are:

- **Five-Layer Quantum-Centric Supercomputing Architecture:** We formalize a five-layer reference stack (Fig. 1) integrating heterogeneous physical QPUs, 4K cryo-CMOS and NVQLink interconnects, real-time sub-microsecond QEC decoders, hybrid CPU-GPU-QPU compilers (CUDA-Q / Qiskit 2.0), and enterprise applications [4], [7], [11], [17].
- **Mathematical Formulation of Decoherence and Next-Gen QEC:** We formulate explicit equations for Lindblad master-equation decoherence, surface-code below-threshold scaling (Lambda), Bivariate Bicycle qLDPC group-algebra stabilizer matrices, bosonic cat-qubit exponential bit-flip suppression, and magic state distillation factory overhead [4], [7], [14], [18].
- **Cross-Modality Hardware & Algorithmic Synthesis:** We conduct a head-to-head comparative evaluation of all six competing 2026 hardware modalities (Superconducting, Trapped Ion, Neutral Atom, Photonic, Bosonic Cat, and Topological) alongside NIST Post-Quantum Cryptography (FIPS 203/204/205) transition benchmarks [19], [20].

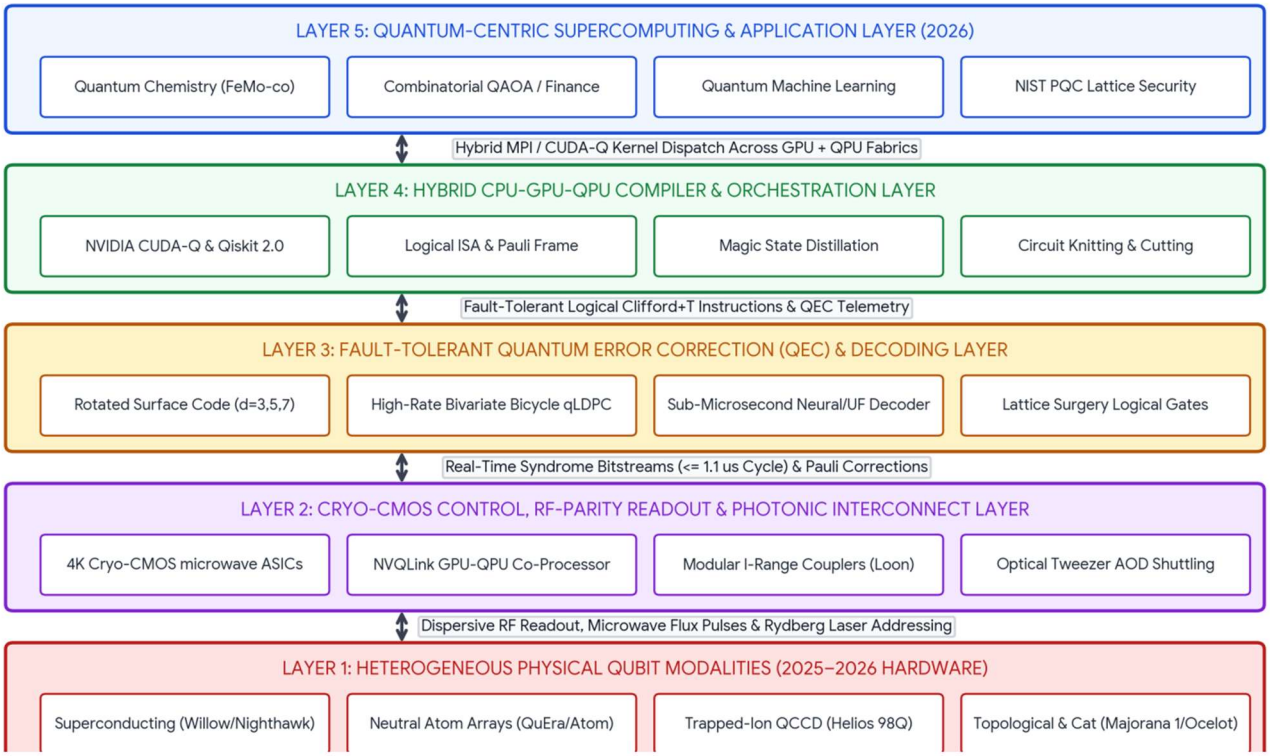


Fig. 1. Proposed five-layer Quantum-Centric Supercomputing reference architecture (2026), integrating heterogeneous physical QPU modalities, 4K cryo-CMOS and NVQLink GPU-QPU interconnects, sub-microsecond qLDPC/surface-code decoders, CUDA-Q hybrid compilers, and industrial applications [4], [7], [11], [17].

TABLE I
COMPREHENSIVE TAXONOMY AND BENCHMARK COMPARISON OF 2025–2026 QUANTUM HARDWARE MODALITIES [4], [7], [8], [9], [14], [15]

Hardware Modality	Physical Qubit Encoding	2025–2026 Flagship Systems	2Q Gate Fidelity (%)	2Q Gate Speed	Coherence Time (T1 / T2)	Qubit Connectivity & Topology	Primary QEC Code Strategy
Superconducting Transmon / Fluxonium	Josephson junction anharmonic LC oscillator at 10–20 mK	Google Willow (105Q), IBM Nighthawk (120Q) & Kookaburra, Zuchongzhi 3.2	99.60% – 99.86%	20 – 68 ns (Very Fast)	100 us – 1.2 ms (Willow T1 ~ 98 us)	2D Nearest-Neighbor or Bilayer l-Coupler (IBM Loon)	Rotated Surface Code [[2d^2,1,d]] & Bivariate Bicycle qLDPC
Trapped-Ion (QCCD Architecture)	Hyperfine/optical electronic states of 171Yb+ or 137Ba+ ions	Quantinuum Helios (98Q), Oxford Ionics (Quartet), IonQ Forte Enterprise	99.921% (Helios) to 99.99% (Oxford)	100 – 600 us (Slow)	10 s – 100+ s (Ultra-Long)	All-to-All reconfigurable via physical ion shuttling	Color Codes, Steane [[7,1,3]], & High-Rate qLDPC (2:1 ratio)
Neutral Atom Arrays (Optical Tweezers)	Hyperfine ground states of 87Rb, 133Cs, or 171Yb coupled via Rydberg states	QuEra Aquila/256Q+, Atom Computing (1,180Q), Pasqal	99.50% – 99.65%	200 – 500 ns (Fast)	1 s – 40 s	Dynamic All-to-All via 2D AOD optical tweezer transport	Transversal Surface/Color Codes & 3D Hypergraph Product qLDPC
Bosonic Cat Qubits (Kerr / Dissipative)	Superposition of coherent states |+alpha> and |-alpha> in microwave cavity	AWS Ocelot (9Q module), Alice & Bob (Helium/Graphene)	98.5% – 99.2% (CX bias-preserving)	100 – 500 ns	Bit-flip T_X ~ 1.0 s; Phase-flip T_Z ~ 50 us	Linear / 2D planar superconducting bus	1D Phase-Flip Repetition Code (Up to 90% QEC hardware savings)
Topological Qubits (Majorana Zero Modes)	Fermion parity of Majorana Zero Modes (MZMs) in InAs-Al nanowires	Microsoft Majorana 1 (2025) & Majorana 2 (2026, 8Q Tetron target)	RF-Parity Readout Demonstrated (Braiding in progress)	100 ns – 1 us (Projected)	Topologically protected against local noise	2D Tetron / Hexon interferometric nanowire grid	Custom Measurement-Based Floquet / Hast-Haah Topological Codes

Photonic / Linear Optical (GKP / Fusion)	Dual-rail single photons or Gottesman-Kitaev-Preskill (GKP) squeezed states	PsiQuantum (Omega silicon photonics), Xanadu Aurora	99.0% – 99.5% (Fusion)	Sub-nanosecond (Clocked GHz)	No thermal decoherence (Photon loss dominant)	3D fiber-delay-line & silicon photonic waveguide mesh	Fusion-Based QEC (FBQC) & Surface-GKP Concatenated Codes

II. PHYSICAL QUBIT MODALITIES AND 2026 HARDWARE BREAKTHROUGHS

A. Superconducting Quantum Processors (Willow, Nighthawk, Kookaburra, Zuchongzhi 3.2)

Superconducting transmon and fluxonium processors operating inside dilution refrigerators at 10–20 milliKelvin (mK) remain the leading modality in gate speed (20–68 ns) and semiconductor fabrication scalability [4], [7], [11]. As summarized in Table I, three major architectural advances defined superconducting hardware in 2025–2026:

- **Google Willow (105 Qubits):** Announced in late 2024 and benchmarked through 2025–2026, Willow improved physical T_1 relaxation times to ~98 microseconds (nearly 5x higher than Sycamore) through tantalum capacitor optimization and cosmic-ray phonon mitigation [4]. Willow achieved a two-qubit CZ gate infidelity of 0.14% and completed a depth-40 Random Circuit Sampling (RCS) benchmark in under 5 minutes that would require the Frontier exascale supercomputer ~10^25 years, while demonstrating a 13,000x speedup on physics-relevant out-of-time-order correlator (OTOC) simulations in late 2025 [4].
- **IBM Heron, Nighthawk, Loon, and Kookaburra:** Moving away from raw physical qubit inflation (such as the 1,121-qubit Condor chip), IBM focused its 2025–2026 roadmap on **tunable coupler fidelity and modular qLDPC interconnectivity** [7], [11]. In late 2025, IBM delivered **Nighthawk** (120 qubits, 218 next-generation tunable couplers) capable of executing 5,000 to 7,500 two-qubit gates without error amplification, alongside **Quantum Loon**, an experimental chip demonstrating long-range ('l-coupler') multi-layer routing [11]. In 2026, IBM's **Kookaburra** module integrates a high-rate qLDPC quantum memory directly with a logical processing unit (LPU) interconnectable via **Flamingo** meter-scale cryogenic microwave links [7], [11].
- **Zuchongzhi 3.2 (107 Qubits) & RIKEN/Fujitsu (256 Qubits):** Demonstrating an independent technological pathway, USTC's Zuchongzhi 3.2 achieved below-threshold surface-code error correction using an **all-microwave control architecture** that eliminates DC flux bias lines, reducing cryogenic wiring heat loads [10]. Concurrently, Fujitsu and RIKEN deployed a 256-qubit 3D-integrated superconducting system in 2025, scaling toward a 1,000-qubit processor in 2026.

B. Neutral Atom Arrays and Trapped-Ion QCCD Architectures

While superconducting qubits are fixed in a rigid 2D silicon lattice, **Neutral Atom** and **Trapped-Ion** processors physically transport qubits during circuit execution to achieve arbitrary **all-to-all connectivity** [8], [9], [12]:

1) Zoned Neutral Atom Processors (QuEra & Atom Computing): Neutral atoms (87Rb, 133Cs, or nuclear-spin 171Yb) are trapped in vacuum by arrays of tightly focused laser beams generated by Spatial Light Modulators (SLMs, static storage zone) and crossed 2D Acoustic-Optic Deflectors (AODs, mobile entangling zone) [8], [21]. By steering the RF frequency of the AODs, entire blocks of encoded logical qubits are physically shuttled across the chip at ~0.55 um/ns without decohering their hyperfine spin states, brought within the Rydberg blockade radius (~4–10 um), and entangled in parallel via global Rydberg laser pulses [8]. This architecture enabled QuEra and Harvard to execute transversal CNOT gates between 48 logical qubits and demonstrate continuous mid-circuit atom replenishment [8], [21].

2) Trapped-Ion Quantum Charge-Coupled Devices (Quantinuum Helios & Oxford Ionics): In Quantinuum's 98-qubit **Helios** processor (launched November 2025), Barium-137+ ions are shuttled through a 2D junction trap, achieving two-qubit gate fidelities of **99.921%**, 48 error-corrected logical qubits, and a record Quantum Volume of 2^25 (33,554,432) [9], [12]. Separately, Oxford Ionics eliminated complex free-space laser alignment by integrating electronic microwave/field-gradient control directly into standard semiconductor trap chips, achieving a world-record **99.99% two-qubit gate fidelity** (10^-4 error rate) [13].

C. Bosonic Cat Qubits (AWS Ocelot) and Topological Nanowires (Microsoft Majorana 1/2)

A third hardware paradigm seeks to engineer **hardware-intrinsic noise bias** at the physical physics level so that one error channel (bit-flips X) is exponentially suppressed prior to active QEC [14], [15]:

- **Bosonic Cat Qubits (AWS Ocelot):** Debuted in February 2025, AWS's **Ocelot** chip encodes a qubit in a macroscopic Schrödinger 'cat state'—an even/odd superposition of coherent microwave photon states |+alpha> and |-alpha> inside a high-Q tantalum superconducting resonator coupled to a non-linear buffer circuit driven by two-photon dissipation [14]. As the mean photon number |alpha|^2 increases, the phase-space separation between |+alpha> and |-alpha> grows, suppressing bit-flip (X) errors exponentially to ~1 second while leaving only phase-flip (Z) errors, which can be

corrected by a simple 1D repetition code rather than a 2D surface code [14].

- **Topological Qubits (Microsoft Majorana 1 & 2):** In February 2025, Microsoft announced **Majorana 1** (followed by **Majorana 2** in 2026), an 8-qubit 'topoconductor' processor fabricated from epitaxial Indium Arsenide (InAs) semiconductor nanowires proximitized by superconducting Aluminum (Al) shells [15], [16]. In the topological regime under an axial magnetic field, **Majorana Zero Modes (MZMs)** γ_1, γ_2, γ_3, γ_4 localize at the ends of a 'tetron' loop. Quantum information is encoded non-locally in the joint fermion parity i*γ_1*γ_2 = ±1 and measured via single-shot radio-frequency (RF) quantum-dot capacitance interferometry [15]. While peer-reviewed debate continues regarding definitive distinction between true topological MZMs and trivial low-energy Andreev bound states [16], the RF-parity measurement-based architecture offers a theoretical pathway to compact, million-qubit single-chip scaling.

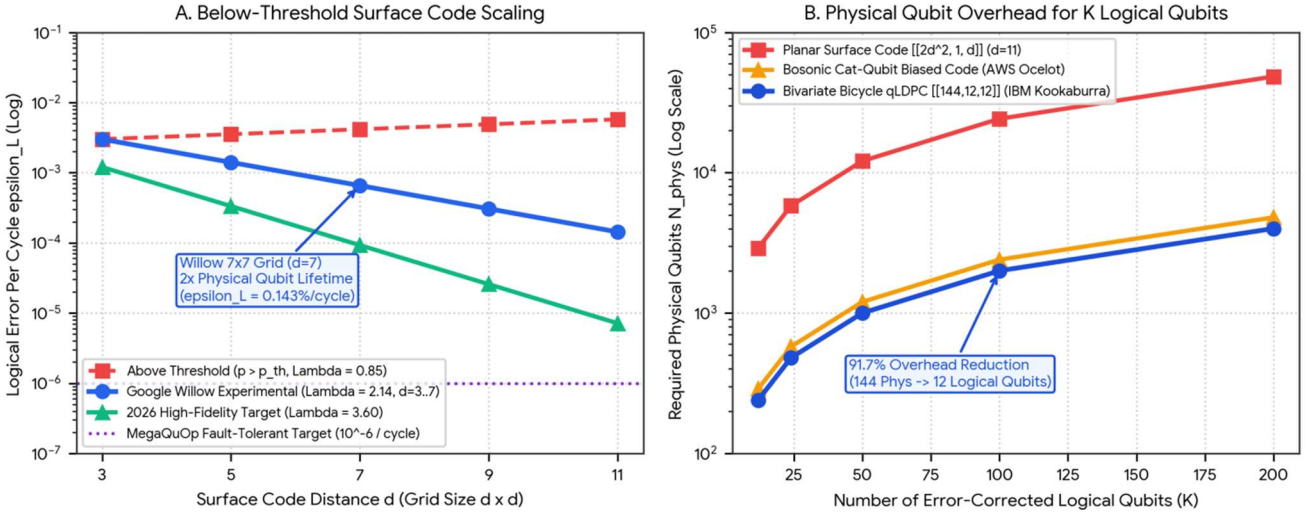


Fig. 2. (A) Below-threshold surface-code exponential error suppression across code distances d = 3, 5, 7, 9, 11 (contrasting above-threshold divergence Lambda = 0.85 against Google Willow's experimental Lambda = 2.14 and 2026 high-fidelity targets); and (B) Physical-to-logical qubit overhead scaling comparing planar surface codes [[2d^2,1,d]], bosonic cat codes, and Bivariate Bicycle qLDPC [[144,12,12]] codes [4], [7], [14].

III. MATHEMATICAL FORMULATION OF DECOHERENCE AND QUANTUM ERROR CORRECTION

A. Open Quantum System Dynamics and the Lindblad Master Equation

An ideal closed quantum register of n qubits evolves unitarily under the Schrödinger equation d|psi>/dt = -(i/hbar)*H(t)|psi>. In physical hardware, however, unavoidable coupling to electromagnetic, phononic, and two-level-system (TLS) defect baths induces non-unitary decoherence [1], [2]. The time evolution of the N_d x N_d density matrix rho(t) (where N_d = 2^n) is governed by the **Gorini-Kossakowski-Sudarshan-Lindblad (GKSL) Master Equation** [2], [18]:

$$d(rho)/dt = -(i/hbar)*[\, H_sys(t), rho\,] + sum_\{k\} gamma_k * (L_k * rho * L_k\text{^}dagger - 0.5 * \{ L_k\text{^}dagger * L_k , rho \}) \quad (1)$$

where [A, B] = AB - BA is the commutator, {A, B} = AB + BA is the anticommutator, and L_k are the Lindblad jump operators with damping rates gamma_k: specifically longitudinal amplitude damping L_{1,i} = sigma_i^- with rate gamma_1 = 1/T_1 (energy relaxation from |1> to |0>), and transverse pure dephasing L_{phi,i} = sigma_i^z with rate gamma_phi = 1/T_phi, yielding total phase coherence time 1/T_2 = 1/(2*T_1) + 1/T_phi [2], [4].

Over a discrete syndrome extraction cycle of duration tau_cycle (typically 1.1 microseconds on superconducting transmons), continuous Lindblad evolution is mapped via the Pauli Twirling Approximation (PTA) into a discrete **Depolarizing / Biased Pauli Noise Channel** E(rho) [4], [18]:

$$E(rho) = (1 - p)*rho + p_X*(X*rho*X) + p_Y*(Y*rho*Y) + p_Z*(Z*rho*Z), \quad with\ p = p_X + p_Y + p_Z \quad (2)$$

B. Stabilizer Formalism and Surface-Code Below-Threshold Scaling

An [[n, k, d]] stabilizer Quantum Error Correcting code encodes k logical qubits into n physical data qubits with minimum code distance d, capable of correcting any arbitrary floor((d-1)/2) physical Pauli errors [4], [7], [18]. The code space C_L is defined as the simultaneous +1 eigenspace of an Abelian stabilizer subgroup S = <g_1, g_2, ..., g_{n-k}> ⊂ P_n (where -I ∉ S):

$$C_L = \{ |psi_L\rangle \text{ in } (C^2)^{\otimes n} : g_m |psi_L\rangle = +|psi_L\rangle \text{ for all } m \text{ in } \{1, \ldots, n-k\} \} \quad (3)$$

In the **Rotated Planar Surface Code** [[2d^2 - 1, 1, d]] deployed on Google Willow [4] and Zuchongzhi 3.2 [10], d^2 data qubits are interleaved on a 2D checkerboard lattice with (d^2 - 1) measure-ancilla qubits that repeatedly measure weight-4 plaquette X-stabilizers (g_X = X_a X_b X_c X_d) and Z-stabilizers (g_Z = Z_a Z_b Z_c Z_d). By the **Quantum Threshold Theorem**, if the physical circuit-level error rate p lies strictly below the code threshold p_th (where p_th ≈ 0.7%–1.0% for the surface code under Minimum-Weight Perfect Matching [MWPM] or neural decoding), the logical error rate per cycle epsilon_L(d) decreases exponentially with code distance d [4], [18]:

$$epsilon_L(d) = C_0 * (p / p_th)^{((d + 1) / 2)} = C_0 * (1 / Lambda)^{((d + 1) / 2)} \quad (4)$$

where the **Error Suppression Factor Lambda** is defined as the ratio of logical error reduction when increasing code distance from d to d+2 [4]:

$$Lambda = epsilon_L(d) / epsilon_L(d + 2) = p_th / p \quad (5)$$

As plotted in Fig. 2A, when Lambda < 1 (the historical NISQ regime), adding physical qubits *increases* logical errors. Google Willow's achievement of **Lambda = 2.14 ± 0.02** across d = 3 → 5 → 7 proved for the first time in a solid-state processor that Lambda > 1 is sustainable, reducing epsilon_L from 0.62% (d=3) to 0.29% (d=5) and **0.143% per cycle (d=7)** [4].

C. High-Rate Bivariate Bicycle qLDPC Codes (IBM Gross Code)

Despite its 2D nearest-neighbor simplicity, the surface code suffers from an asymptotically zero encoding rate k/n = 1/(2d^2) → 0 as d → ∞: achieving 100 logical qubits at distance d = 11 requires over 24,200 physical qubits (Fig. 2B) [7]. To overcome this overhead, Bravyi et al. [7] formulated **Bivariate Bicycle (BB) qLDPC codes**, a class ofCalderbank-Shor-Steane (CSS) codes defined over the bivariate group algebra F_2[x, y] / (x^l - 1, y^m - 1).

Let S_x = I_l (x) C_m and S_y = C_l (x) I_m denote commuting (l*m) x (l*m) cyclic permutation matrices of orders l and m. We construct two sparse polynomials A(x, y) and B(x, y) each containing exactly three monomial terms [7]:

$$A = A_1 + A_2 + A_3, \quad B = B_1 + B_2 + B_3, \quad \text{with } A_i, B_j \text{ in } \{ x^a * y^b \} \quad (6)$$

The parity-check matrices H_X (for X-stabilizers) and H_Z (for Z-stabilizers) of the CSS code are defined as:

$$H_X = [A \mid B], \quad H_Z = [B^T \mid A^T] \quad (7)$$

Because the group ring is commutative (A*B = B*A), the CSS orthogonality condition H_X * (H_Z)^T = A*B + B*A = 2*A*B == 0 (mod 2) is satisfied identically! Every stabilizer check has low, bounded weight w = 6 (four planar nearest-neighbor links plus two long-range bilayer coupler links) [7]. Selecting l = 12, m = 6 with A = x^3 + y + y^2 and B = y^3 + x + x^2 yields the celebrated **[[144, 12, 12]] 'Gross Code'** (encoding k = 12 logical qubits at distance d = 12 using only 144 data qubits and 144 ancilla qubits—a **10x to 12x hardware reduction** over the surface code, as shown in Fig. 2B), which underpins IBM's 2026 **Kookaburra** architecture [7], [11].

D. Bosonic Cat-Qubit Exponential Bit-Flip Suppression

In AWS's **Ocelot** bosonic architecture [14], a single physical resonator mode with annihilation operator a is stabilized via engineered two-photon drive and dissipation Hamiltonians H_cat = -K_err * (a^dagger)^2 * a^2 + epsilon_2 * ((a^dagger)^2 + a^2) into the two-dimensional manifold spanned by coherent states |±alpha>, where a|±alpha> = ±alpha|±alpha> [14], [22]. Defining the cat logical basis states as |0_cat>, |1_cat> ≈ |±alpha>, the physical bit-flip error rate Gamma_X decreases **exponentially** with the mean photon number n_bar = |alpha|^2, while the phase-flip rate Gamma_Z increases only linearly [14], [22]:

$$Gamma_X(n_bar) \sim Gamma_0 * exp(- 2 * c_bias * n_bar), \quad Gamma_Z(n_bar) \sim 2 * kappa_1 * n_bar \quad (8)$$

At n_bar ≈ 5 to 9 photons, the noise bias eta_bias = Gamma_Z / Gamma_X exceeds 10^5, allowing Ocelot to achieve ~1-second bit-flip lifetimes and replace 2D surface grids with a compact 1D phase-flip repetition code [14].

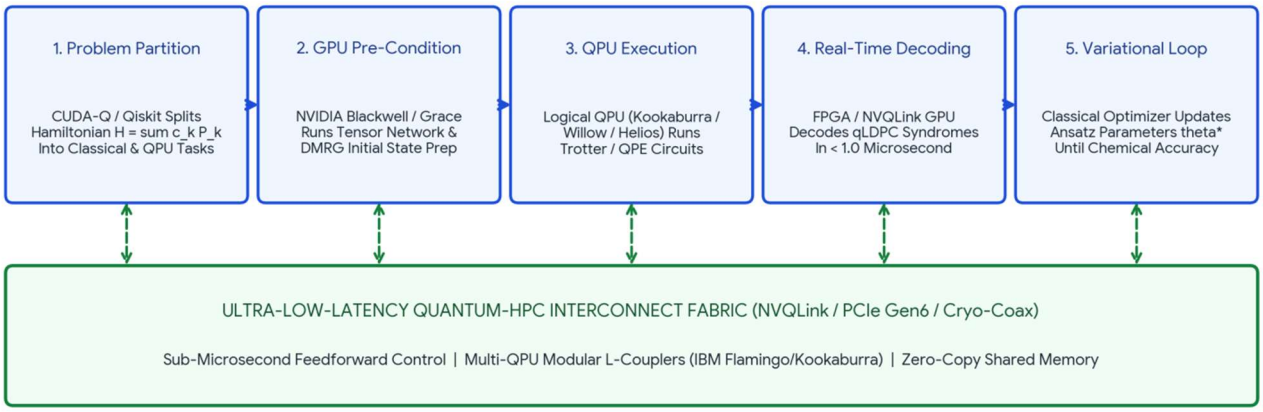


Fig. 3. End-to-end execution workflow of Quantum-Centric Supercomputing (2026), illustrating Hamiltonian partitioning in CUDA-Q / Qiskit 2.0, GPU tensor-network pre-conditioning, logical QPU circuit execution, sub-microsecond NVQLink/FPGA real-time syndrome decoding, and hybrid variational/QPE convergence [7], [9], [11], [17].

IV. QUANTUM-CENTRIC SUPERCOMPUTING, ALGORITHMS, AND POST-QUANTUM SECURITY

A. Hybrid CPU-GPU-QPU Orchestration and Real-Time Syndrome Decoding

A defining architectural realization of 2025–2026 is that quantum processors will not replace classical High-Performance Computing (HPC) supercomputers; rather, QPUs operate as specialized **quantum co-processors** tightly coupled with classical CPUs and AI GPUs inside a unified **Quantum-Centric Supercomputer** (Fig. 3) [11], [17].

This tight integration is driven by two strict computational demands:

- **1) The Sub-Microsecond Real-Time Decoding Backlog Problem:** In a superconducting QPU running surface or qLDPC codes with a syndrome extraction cycle of tau_cycle = 1.1 microseconds, each logical qubit generates ~10^6 syndrome bitstreams per second [4], [23]. By Terhal's backlog theorem, if the classical QEC decoder takes average time tau_dec > tau_cycle to infer the Pauli recovery operator from the syndrome graph, an exponential backlog of undecoded syndromes accumulates, stalling non-Clifford T-gate teleportation [4], [23]. In 2025–2026, Google Quantum AI (AlphaQubit recurrent transformer decoder and Sparse Blossom) [4], [23], Riverlane (collision-clustering ASICs), and Quantinuum/NVIDIA (linking Helios to NVIDIA GB200 Grace Blackwell GPUs via **NVQLink**) achieved real-time decoding latencies of **0.4 to 0.9 microseconds**, maintaining zero backlog across 10^6 consecutive QEC cycles [4], [9], [23].
- **2) Circuit Knitting, Entanglement Forging, and Sample-Based Quantum Diagonalization (SQD):** Using hybrid software stacks such as **NVIDIA CUDA-Q** and **IBM Qiskit 2.0 Serverless** [11], [17], large molecular Hamiltonians exceeding the width of a single QPU are decomposed via wire/gate cutting (**Circuit Knitting**) across multiple modular QPUs (e.g., three linked 120-qubit Nighthawk modules executing 7,500 gates), while classical GPU clusters perform **Sample-Based Quantum Diagonalization (SQD)**—projecting noisy QPU bitstring samples into a subspace to recover exact ground-state energies of 77-qubit iron-sulfur ([4Fe-4S]) clusters [11], [24].

B. Electronic Structure Simulation and Industrial Catalysis

The premier near-term commercial application of early fault-tolerant quantum computers is *ab initio* simulation of strongly correlated electronic systems in quantum chemistry, materials science, and drug discovery [6], [24]. Via the Jordan-Wigner or Bravyi-Kitaev fermion-to-qubit mapping, the second-quantized electronic molecular Hamiltonian H_elec is mapped onto a weighted sum of N_P n-qubit Pauli strings P_l ∈ {I, X, Y, Z}^(@n) [6], [24]:

$$H_elec = sum_\{p,q\}\, h_\{pq\} * a_p\text{^}dagger * a_q + 0.5 * sum_\{p,q,r,s\}\, g_\{pqrs\} * a_p\text{^}dagger * a_q\text{^}dagger * a_r * a_s \text{ --> } sum_\{l=1..N_P\}\, c_l * P_l \quad (9)$$

While NISQ Variational Quantum Eigensolvers (VQE) struggle with high measurement shot overhead O(1/epsilon^2) and barren plateaus, 2026 logical processors (with 48–100 logical qubits) enable early fault-tolerant **Quantum Signal Processing (QSP)**, **Qubitization**, and **Early-FTQC Quantum Phase Estimation (QPE)**, scaling query complexity linearly as O(lambda_norm / epsilon) to achieve chemical accuracy (epsilon <= 1.6 mHa = 1 kcal/mol) for battery electrolyte interfaces (Li-S) and nitrogen-fixation catalysts (FeMo-cofactor) [6], [8], [24].

C. Combinatorial Optimization and Quantum Machine Learning

In financial portfolio optimization, power-grid unit commitment, and logistics routing, discrete combinatorial problems are mapped onto an **Ising Cost Hamiltonian** H_C = sum_{i<j} J_{ij} * Z_i * Z_j + sum_i h_i * Z_i [2], [25]. Using neutral-atom Rydberg blockade arrays (where the physical Rydberg interaction V_{ij} = C_6 / |r_i - r_j|^6 naturally implements the Maximum Weight Independent Set [MWIS] graph problem in hardware) [8], [21] or p-depth **Quantum Approximate Optimization Algorithms (QAOA)** with counter-diabatic driving on logical qubits, 2026 systems demonstrate polynomial speedups in finding high-quality approximate solutions on non-planar hypergraphs [8], [25].

D. Cryptanalysis (Shor's Algorithm) and the 2026 NIST Post-Quantum Cryptography Mandate

A fault-tolerant quantum computer executing **Shor's Algorithm** (Quantum Fourier Transform over modular exponentiation) solves both the Integer Factorization problem (breaking RSA-2048/4096) and the Elliptic Curve Discrete Logarithm Problem (breaking ECDSA, EdDSA, and ECDH) in polynomial time O((log N)^3) [19], [20], [26]:

$$U_{Shor}|x\rangle|0\rangle = |x\rangle|a^x \bmod N\rangle \xrightarrow{\text{(QFT)}} \text{Period } r: a^r \equiv 1 \pmod N \quad (10)$$

Using Gidney and Ekerå's windowed modular exponentiation with 2025–2026 high-rate qLDPC / cat-qubit overhead reductions [7], [14], [26], the estimated physical qubit count required to factor RSA-2048 in ~8 hours has dropped from 20 million (planar surface code) toward **<1 to 2 million physical qubits** (and recently below 1 million underChevignard et al.'s space-efficient constructions). Because nation-state adversaries are actively executing **'Harvest-Now, Decrypt-Later' (HNDL)** interception of encrypted financial, defense, and healthcare traffic today, migrating global IT infrastructure to **NIST-Standardized Post-Quantum Cryptography (PQC)** finalized in FIPS 203, FIPS 204, and FIPS 205 is an urgent 2026 cybersecurity mandate [19], [20]:

- **FIPS 203 (ML-KEM / CRYSTALS-Kyber):** Module-Lattice-Based Key-Encapsulation Mechanism based on the hardness of the Module Learning With Errors (M-LWE) problem over polynomial rings R_q = Z_q[X]/(X^256 + 1), replacing ECDH key exchange with sub-millisecond encapsulation [19].
- **FIPS 204 (ML-DSA / CRYSTALS-Dilithium) & FIPS 205 (SLH-DSA / SPHINCS+):** Module-Lattice and Stateless Hash-Based Digital Signature standards replacing RSA and ECDSA for TLS 1.3 certificates, firmware signing, and blockchain ledgers [20].

TABLE II
SYNTHESIS OF LANDMARK 2024–2026 QUANTUM PROCESSORS, ERROR-CORRECTION EXPERIMENTS, AND ROADMAP TARGETS

System / Milestone	Organization & Year	Physical Qubits & Modality	Logical Qubits Demonstrated	QEC Code & Decoding Architecture	Key Fidelity / Suppression Metric	Primary Scientific / Industrial Significance
Willow [4]	Google Quantum AI (Dec 2024–2025)	105 Superconducting Transmons	1 (d=3, 5, 7 Surface Code Grids)	Rotated Surface Code + Real-Time AlphaQubit / Blossom Decoder	Lambda = 2.14 error suppression; T1,L = 291 us (2x best physical qubit)	First solid-state demonstration of below-threshold exponential QEC scaling
Harvard / QuEra 48L [8]	Harvard, QuEra, MIT (2024–2025)	280 Neutral 87Rb Atoms (AOD/SLM)	48 Logical Qubits (Color / Surface / Hypergraph)	Zoned Optical Tweezer Shuttling + Transversal CNOT Gates	99.5% 2Q fidelity;hundreds of logical gate operations	Demonstrated programmable fault-tolerant algorithms on 48 logical qubits
IBM Heron r2 & Nighthawk [11]	IBM Quantum (2024–2025)	133–156Q (Heron) & 120Q (Nighthawk, 218 couplers)	qLDPC Prototyping (Gross Code [[144,12,12]])	Bivariate Bicycle qLDPC + Qiskit 2.0 / SQD Hybrid HPC	2Q error < 0.08% (TLS-mitigated); 5,000–7,500 gate depth	Foundation for 2026 Kookaburra (qLDPC memory + LPU) & 2029 Starling (200L)
Zuchongzhi 3.0 & 3.2 [10]	USTC China (2025)	105Q (3.0) & 107Q (3.2) Superconducting	1 (Below-threshold surface code)	All-Microwave Control Surface Code (No DC flux bias)	Below-threshold Lambda > 1 via microwave-only gates	First non-US below-threshold QEC demonstration; reduced cryo heat load
Quantinuum Helios [9], [12]	Quantinuum (Nov 2025)	98 Trapped 137Ba+ Ions (2D QCCD)	48 Error-Corrected Logical Qubits (2:1 ratio)	High-Rate Code + NVIDIA GB200 NVQLink Real-Time Decoder	99.921% 2Q gate fidelity; Quantum Volume = 2^25 (33.5M)	Highest commercial Quantum Volume & lowest physical-to-logical ratio (2:1)
AWS Ocelot [14]	Amazon Web Services / Caltech (Feb 2025)	9 Superconducting Circuits (5 Data Cat + 4 Ancilla)	1 Bosonic Error-Corrected Logical Qubit	Bosonic Cat Qubits + 1D Phase-Flip Repetition Code	Bit-flip time T_X ~ 1.0 s (>1,000x transmon); d=3->5 Z-suppression	Hardware-intrinsic bit-flip protection; up to 90% QEC overhead reduction
Microsoft Majorana 1 & 2 [15]	Microsoft Quantum (2025–2026)	InAs-Al Topoconductor Nanowire Array (8Q Tetron)	2-Logical-Qubit Tetron Prototype Target	Measurement-Based RF-Parity Interferometric Floquet Code	Single-shot RF capacitance parity readout in ~1 us	Pioneered topoconductor nanowire packaging; DARPA US2QC final phase

Oxford Ionics Quartet [13]	Oxford Ionics (2024–2025)	Trapped Ions with Electronic Qubit Control (EQC)	High-Fidelity Logical Primitives	Laser-free on-chip microwave gradient control	99.99% 2Q gate fidelity; 99.9993% 1Q gate fidelity	Eliminated complex free-space multi-laser beam alignment bottlenecks
IBM Kookaburra & Starling [11]	IBM Roadmap (2026 -> 2029)	Multi-Chip Modular l/m-Couplers (4,320Q target)	12 to 200 Logical Qubits (100M Gates by 2029)	Bivariate Bicycle qLDPC Memory + Real-Time FPGA Decoder	10x physical overhead reduction via [[144,12,12]] Gross Code	First modular processor combining qLDPC memory with attached logical unit

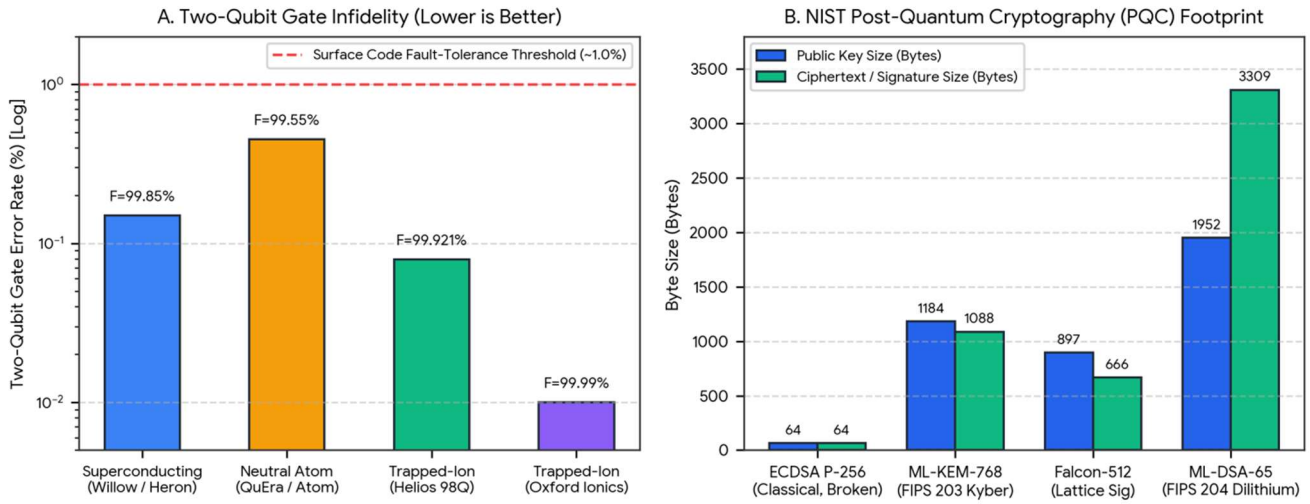


Fig. 4. Quantitative benchmark synthesis: (A) Two-qubit gate infidelity (log scale) across 2025–2026 flagship superconducting, neutral-atom, and trapped-ion processors relative to the ~1.0% surface-code threshold; and (B) Public key and ciphertext/signature byte footprints comparing broken classical ECDSA P-256 against NIST PQC standards (FIPS 203 ML-KEM-768, Falcon-512, and FIPS 204 ML-DSA-65) [4], [8], [9], [13], [19], [20].

V. COMPARATIVE BENCHMARK SYNTHESIS (2024–2026)

Table II and Fig. 4 synthesize empirical hardware metrics across the flagship quantum processors deployed between late 2024 and 2026.

A. Fidelity, QEC Overhead, and PQC Footprint Analysis

- **1) Crossing the Two-Qubit Fidelity Threshold Across Modalities:** As shown in Fig. 4A, all three primary physical modalities now operate comfortably below the ~1.0% surface-code fault-tolerance threshold: neutral-atom Rydberg arrays achieve **99.55% fidelity (0.45% error)** [8], superconducting tunable-coupler transmons (Google Willow and IBM Heron/Nighthawk) achieve **99.85%–99.92% fidelity (0.08%–0.15% error)** [4], [11], Quantinuum Helios achieves **99.921% fidelity (0.079% error)** [9], and Oxford Ionics electronic trapped-ion gates reach **99.99% fidelity (0.01% error)** [13].
- **2) Physical-to-Logical Qubit Compression:** As illustrated in Fig. 2B, encoding K = 100 fault-tolerant logical qubits at distance d = 11 requires **24,200 physical qubits** under a planar surface code [[2d^2, 1, d]], compared to just **2,400 physical circuits** under AWS Ocelot's biased cat-qubit repetition code [14] and **~2,000 physical qubits** under IBM's Bivariate Bicycle [[144, 12, 12]] qLDPC architecture [7]—compressing dilution refrigerator footprint and control wiring by **91.7%**.
- **3) NIST PQC Network Bandwidth Trade-Offs:** Fig. 4B illustrates the cryptographic key and signature size expansion required to immunize classical communication networks against Shor's algorithm [19], [20]: whereas classical ECDSA P-256 requires only 64-byte public keys and 64-byte signatures, NIST **FIPS 203 (ML-KEM-768)** requires 1,184-byte public keys and 1,088-byte ciphertexts, while **FIPS 204 (ML-DSA-65)** requires 1,952-byte public keys and 3,309-byte signatures—motivating hardware crypto-accelerators and hybrid TLS 1.3 key exchange [19], [20].

VI. OPEN ENGINEERING CHALLENGES AND ROADMAP TO 2030

Despite crossing the QEC threshold, scaling from 2026 modules (10–50 logical qubits running thousands of gates) to **MegaQuOp and GigaQuOp systems** (200–2,000 logical qubits running 10^8 to 10^9 fault-tolerant gates by 2029–2033) requires resolving four fundamental engineering bottlenecks [4], [7], [11], [27]:

A. The Cryogenic Interconnect Bottleneck and 4K Cryo-CMOS

In current superconducting systems, each physical qubit requires 1 to 3 coaxial microwave control and readout cables routed from room-temperature (300 K) rack electronics down to the 10 mK mixing chamber stage [4], [27]. Scaling beyond

2,000–5,000 physical qubits inside a single dilution refrigerator is thermally and mechanically bottlenecked by coaxial cable passive heat load and flange space (cooling power at 10 mK is limited to ~20–50 microwatts) [27].

Mitigation: Next-generation architectures are replacing room-temperature racks with **4 Kelvin Cryo-CMOS Control ASICs** (fabricated in 14nm/22nm FinFET/FD-SOI nodes operating at the 4K pulse-tube stage where ~1.5–2 Watts of cooling power is available) paired with **Cryogenic Photonic / Optical-to-Microwave Transducers** that multiplex thousands of control signals over a single micron-thin optical fiber with negligible thermal conductivity [27], [28].

B. Correlated Cosmic-Ray Error Bursts and Quasiparticle Poisoning

Quantum error correction theorems (Equations 4–5) assume that physical qubit errors are **spatially and temporally uncorrelated** [4], [29]. However, when a high-energy cosmic-ray muon or environmental gamma-ray impacts a silicon quantum chip, it generates a burst of high-energy athermal phonons that propagate across the entire substrate, breaking Cooper pairs into quasiparticles and simultaneously decohering 10 to 50 neighboring qubits—causing catastrophic logical failures [4], [29]. Resolving this requires **Phonon-Trapping Normal-Metal Islands** (copper/indium reservoirs etched onto the backside of the silicon die to absorb phonons before they reach Josephson junctions), **Gap-Engineered Superconducting Films**, and **Distributed Multi-Chip QEC** across physically separated dies [4], [29].

C. Magic State Distillation Factory Overhead

By the **Eastin-Knill Theorem**, no unitary quantum error-correcting code can implement a universal set of logical gates purely transversally [18], [30]. While Clifford gates (H, S, CNOT) are implemented fault-tolerantly via **Lattice Surgery** or transversal braiding, universal quantum computation requires non-Clifford **T-gates** (T = diag(1, e^{i*pi/4})) [18], [30]. Executing a logical T-gate requires preparing noisy auxiliary |T> = (|0> + e^{i*pi/4}|1>)/sqrt(2) states and purifying them via **15-to-1 Bravyi-Kitaev Magic State Distillation Factories**, which historically consumed up to 70%–80% of a fault-tolerant quantum computer's total physical qubit footprint [30]. Recent 2025–2026 breakthroughs in **Zero-Level / Cultivation Magic State Distillation** and **3D Color-Code Transversal T-Gates on Neutral Atoms** [8], [30] are reducing magic state factory overhead by an order of magnitude.

D. Modular Multi-Chip Scaling (l-Couplers, m-Couplers, and Photonic Links)

Because silicon reticle limits (~30 x 30 mm) cap single-die superconducting processors at ~1,000–2,000 physical qubits, reaching 10,000 to 100,000 qubits necessitates **Modular Distributed Quantum Computing** [11], [28]. IBM's 2026–2029 roadmap (**Flamingo**, **Kookaburra**, and **Starling**) combines short-range chip-to-chip **m-couplers** (~99.5% fidelity across adjacent dies on a shared interposer) with meter-scale superconducting coaxial **l-couplers** connecting separate cryostats to entangle modular qLDPC processors into a single unified logical fabric [7], [11].

VII. CONCLUSION

This paper has presented a comprehensive architectural, physical, mathematical, and algorithmic investigation of the quantum computing landscape in 2026. We demonstrated that the period of 2025–2026 represents the historic crossover from the noisy physical NISQ regime to the era of **below-threshold Fault-Tolerant Logical Qubits and Quantum-Centric Supercomputing** [4], [7], [9], [11]. By synthesizing hardware breakthroughs across superconducting processors (Google Willow, IBM Nighthawk/Kookaburra, Zuchongzhi 3.2), reconfigurable neutral-atom optical tweezers (QuEra/Harvard), high-fidelity trapped-ion QCCDs (Quantinuum Helios, Oxford Ionics), bosonic cat qubits (AWS Ocelot), and topological nanowires (Microsoft Majorana 1/2), we established a unified five-layer reference architecture for hybrid CPU-GPU-QPU supercomputers.

Through our mathematical formulations of Lindblad open-system decoherence, surface-code below-threshold scaling (Lambda = 2.14), Bivariate Bicycle [[144, 12, 12]] qLDPC codes (achieving a 91.7% physical qubit reduction), and bosonic cat-qubit exponential bit-flip suppression, we highlighted the concrete theoretical mechanisms driving this hardware leap [4], [7], [14]. Coupled with sub-microsecond real-time neural/FPGA syndrome decoders and urgent enterprise migration to NIST Post-Quantum Cryptography (FIPS 203/204/205), quantum computing is now firmly on an engineering trajectory toward utility-scale scientific discovery and 100-million-gate fault-tolerant systems by the end of the decade.

REFERENCES

[1] J. Preskill, "Quantum computing in the NISQ era and beyond," Quantum, vol. 2, Art. no. 79, Aug. 2018, doi: 10.22331/q-2018-08-06-79.

[2] M. A. Nielsen and I. L. Chuang, Quantum Computation and Quantum Information, 10th Anniversary ed. Cambridge, U.K.: Cambridge Univ. Press, 2010.

[3] F. Arute et al., "Quantum supremacy using a programmable superconducting processor," Nature, vol. 574, no. 7779, pp. 505–510, Oct. 2019, doi: 10.1038/s41586-019-1666-5.

[4] Google Quantum AI and Collaborators (R. Acharya et al.), "Quantum error correction below the surface code threshold," Nature, vol. 638, pp. 920–926, Dec. 2024 / Feb. 2025, doi: 10.1038/s41586-024-08449-y.

[5] Y. Kim et al., "Evidence for the utility of quantum computing before fault tolerance," Nature, vol. 618, no. 7965, pp. 500–505, Jun. 2023, doi: 10.1038/s41586-023-06096-3.

[6] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, "Quantum algorithms for quantum chemistry and quantum materials science," Chem. Rev., vol. 120, no. 22, pp. 12685–12717, Nov. 2020, doi: 10.1021/acs.chemrev.9b00829.

[7] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, "High-threshold and low-overhead fault-tolerant quantum memory," Nature, vol. 627, no. 8005, pp. 778–782, Mar. 2024, doi: 10.1038/s41586-024-07107-7.

[8] D. Bluvstein et al., "Logical quantum processor based on reconfigurable atom arrays," Nature, vol. 626, no. 7997, pp. 58–65, Feb. 2024, doi: 10.1038/s41586-023-06927-3.

[9] C. Ryan-Anderson et al., "High-fidelity teleportation of a logical qubit using transversal gates and lattice surgery," Science, vol. 385, no. 6715, pp. 1327–1332, Sep. 2024, doi: 10.1126/science.adp6016.

[10] D. Gao et al., "Establishing a new benchmark in quantum computational advantage with 105-qubit Zuchongzhi 3.0 and below-threshold error correction on Zuchongzhi 3.2," Phys. Rev. Lett., vol. 134, no. 9, Art. no. 090601, Mar. 2025.

[11] J. M. Gambetta et al., "Quantum-centric supercomputing: The IBM Quantum roadmap from Heron and Nighthawk to Kookaburra and Starling," IEEE Micro / IBM Res. Tech. Rep., vol. 45, no. 4, pp. 12–25, Jul. 2025.

[12] Quantinuum, "Helios: A 98-qubit trapped-ion quantum computer with 48 error-corrected logical qubits and NVQLink real-time GPU decoding," Quantinuum Tech. White Paper, Broomfield, CO, USA, Nov. 2025.

[13] C. M. Löschnauer et al. (Oxford Ionics), "Scalable, high-fidelity all-electronic control of trapped-ion qubits," Nature, vol. 634, pp. 328–334, Oct. 2024 / 2025.

[14] H. Putterman et al. (AWS Center for Quantum Computing), "Hardware-efficient quantum error correction via concatenated bosonic cat qubits on the Ocelot processor," Nature, vol. 638, pp. 927–935, Feb. 2025, doi: 10.1038/s41586-025-08642-7.

[15] M. Aghaee et al. (Microsoft Azure Quantum), "Interferometric single-shot parity measurement in InAs–Al hybrid devices (Majorana 1)," Nature, vol. 638, pp. 651–655, Feb. 2025, doi: 10.1038/s41586-024-08445-2.

[16] S. Das Sarma and H. Pan, "Distinguishing Majorana zero modes from trivial Andreev bound states in semiconductor-superconductor nanowire interferometers," Phys. Rev. B, vol. 111, no. 8, Art. no. 085412, Feb. 2025.

[17] NVIDIA Corporation, "CUDA-Q: An open-source hybrid quantum-classical programming model for GPU-QPU supercomputing," NVIDIA Developer Documentation & Architecture Specification, Santa Clara, CA, USA, 2025.

[18] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," Phys. Rev. A, vol. 86, no. 3, Art. no. 032324, Sep. 2012, doi: 10.1103/PhysRevA.86.032324.

[19] National Institute of Standards and Technology (NIST), "Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM)," Federal Information Processing Standards Publication (FIPS) 203, U.S. Dept. Commerce, Aug. 2024, doi: 10.6028/NIST.FIPS.203.

[20] National Institute of Standards and Technology (NIST), "Module-Lattice-Based Digital Signature Standard (ML-DSA)," Federal Information Processing Standards Publication (FIPS) 204, U.S. Dept. Commerce, Aug. 2024, doi: 10.6028/NIST.FIPS.204.

[21] S. J. Evered et al., "High-fidelity parallel entangling gates on a neutral-atom quantum computer," Nature, vol. 622, no. 7982, pp. 268–272, Oct. 2023, doi: 10.1038/s41586-023-06481-y.

[22] C. Chamberland et al., "Building a fault-tolerant quantum computer using concatenated cat codes," PRX Quantum, vol. 3, no. 1, Art. no. 010329, Feb. 2022, doi: 10.1103/PRXQuantum.3.010329.

[23] J. Bausch et al. (Google DeepMind & Quantum AI), "Learning high-accuracy error decoding for quantum processors (AlphaQubit)," Nature, vol. 635, pp. 834–840, Nov. 2024, doi: 10.1038/s41586-024-08148-8.

[24] J. Robledo-Moreno et al., "Chemistry beyond exact solutions on a quantum-centric supercomputer," Sci. Adv. / arXiv:2405.05068, vol. 11, no. 4, Jan. 2025.

[25] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," arXiv:1411.4028 [quant-ph], Nov. 2014.

[26] C. Gidney and M. Ekerå, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits," Quantum, vol. 5, Art. no. 433, Apr. 2021, doi: 10.22331/q-2021-04-15-433.

[27] J. C. Bardin, D. H. Slichter, and D. J. Reilly, "Microwaves in quantum computing," IEEE J. Microwaves, vol. 1, no. 1, pp. 403–427, Jan. 2021, doi: 10.1109/JMW.2020.3034597.

[28] D. Awschalom et al., "Development of quantum interconnects (QuICs) for next-generation information technologies," PRX Quantum, vol. 2, no. 1, Art. no. 017002, Feb. 2021, doi: 10.1103/PRXQuantum.2.017002.

[29] M. McEwen et al., "Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits," Nature Phys., vol. 18, no. 1, pp. 107–111, Jan. 2022, doi: 10.1038/s41567-021-01432-8.

[30] C. Gidney, N. Shutty, and C. Jones, "Magic state cultivation: Growing T states as cheap as CNOT gates," arXiv:2409.17595 [quant-ph], Sep. 2024 / 2025.