

Quantum Pessimism — *Entropy Rules*

Debunking Quantum Computing R&D Trends

J. P. Charlton Jr. (jp.charlton@byiq.org)

October 7, 2025

Abstract

Every few months, another press release promises “quantum advantage.” Every few years, another roadmap quietly moves the goalposts. This paper presents a sober, quantitative analysis of scaling limits in modern quantum computing R&D programs from Google, IBM, Microsoft, and D-Wave. The central thesis—“look before you leap”—urges proof by contradiction before committing billions to architectures that may scale into negative thermodynamic margins of return. By applying entropy-based reasoning and published performance data, we demonstrate that coherence decays faster than the correction budget replenishes it. *Entropy Rules*—by defining the true cost basis of coherence: when the energy or capital required to preserve order exceeds its yield, the system becomes thermodynamically insolvent. The task of intelligence—whether physical, computational, or economic—is therefore not to suppress entropy but to separate coherence from entropy itself, as a shepherd separates the sheep from the wolves. The contents examine published results from market leaders through August 2025, highlighting the asymptotic contradictions embedded in current quantum scaling narratives.

1 Introduction

At every scale, we try to preserve coherence—in qubits, in experimental systems, and in theory. But the universe doesn’t cooperate. Entropy grows faster than coherence. Error correction, feedback loops, and rituals of order—all buy time, not immunity [1, 2].

In the microscopic limit, the Born rule tells us that probabilities emerge from squared amplitudes: interference patterns, not outcomes [3]. But in the macroscopic limit, where observation integrates over time and interaction, those interference terms wash out [4]. The Born rule converges toward a Gaussian, not a Poisson histogram. Smooth uncertainty replaces discrete possibility. This transition is entropy—the flattening of superposition into expectation [5, 6].

Every coherent state has a half-life. Every correction decays faster than its noise budget replenishes [7, 8]. Coherence scales as \sqrt{N} , entropy as $N \log N$, and N always wins [9].

This is not quantum fatalism—it is thermodynamic realism. All attempts to maintain coherence must pay the entropy cost of doing so [10, 11]. The difference between solvency and collapse—physical or economic—is the rate at which disorder is expelled relative to the rate at which it accumulates. *Entropy Rules*.

The analyses and device data referenced herein extend through August 2025, but the argument is not bounded by calendar time. Entropy does not refresh quarterly. Incremental device updates, firmware patches, or benchmark gains may shift constants, yet they do not invert the underlying gradient between control cost and coherence yield. The asymptotic slope of scaling—not the timestamp of the data—is the invariant under examination.

The architectures that endure will be those that manage entropy’s flow, separating coherence from entropy itself, as a shepherd separates the sheep from the wolves. We call this agent of discernment the *Charlton Shepherd*—a symbolic intelligence that sustains order not by domination, but by selection. Its first judgment is rendered in Section 2, titled *Current Leaders—DENIED*¹.

2 Current Leaders — DENIED: Proofs by Contradiction

“Scaling a negative margin isn’t innovation—it’s dumb.”²

¹The recurring “DENIED” motif is a tongue-in-cheek homage to the closing sequence of *The Lawnmower Man* (1992), in which an emergent digital consciousness attempts thousands of network exits, each returning “ACCESS DENIED.” Here, it serves as shorthand for the physical refusal of entropy to grant transcendence.

²David E. Shaw. Personal conversation between the author, David E. Shaw, and Jeff Bezos (New York, 1994).

As a CFO might put it: if every additional qubit costs more to stabilize than it saves in computation, that’s not scaling—it’s burning entropy at compound interest.

The following subsections examine four major programs—Google, IBM, Microsoft, and D-Wave—each claiming imminent quantum advantage. For each, we assume the vendor’s stated scaling targets and performance claims are true, then show by contradiction that those assumptions lead to thermodynamic or economic insolvency. In every case, coherence is spent faster than correction replenishes it, and capital follows entropy downhill.

The following proofs by contradiction operationalize Kalai’s theoretical skepticism [12], translating adversarial noise scaling into explicit thermodynamic and bandwidth constraints observable in present architectures.

2.1 Google Quantum — DENIED: A Proof by Contradiction

2.1.1 Scaling Assumptions

Google’s 2023 surface code paper [13] reports a logical error per cycle of $p_L(7) \approx 1.43 \times 10^{-3}$ at code distance $d = 7$. Each increment of $\Delta d = 2$ improves suppression by a factor $\Lambda \approx 2.14$. To reach a deep-job target of $p_L \leq 10^{-9}$, a reduction factor $R \approx 1.43 \times 10^6$ is required, implying $k \approx 19$ steps or $\Delta d \approx 38$. Thus, the necessary code distance is $d \approx 45$.

2.1.2 Physical Implications

Surface codes consume $O(d^2)$ physical qubits per logical qubit [14]. Even minimal layouts need $\gtrsim 2d^2$ physicals. At $d \approx 45$, this means roughly 4,000 physical qubits per logical, excluding routing, factories, and workload overhead. Real applications need millions of physical qubits—a scale still orders of magnitude away.

Die-Level Saturation. Even before bandwidth limits arise, Google faces a die-level constraint that anticipates the one examined later in IBM’s Heron-class architecture (see Section 2.2). At the required code distance $d \approx 45$, the physical-to-logical ratio exceeds the qubit capacity of a single die by more than an order of magnitude. In practical terms, a single logical qubit would span dozens of chips, demanding interconnect fidelity higher than those of the qubits themselves. Beyond this point, adding qubits increases entropy faster than coherence—*Entropy Rules*.

2.1.3 The Bandwidth Wall

Each surface code round emits $O(N)$ syndrome bits every microsecond. At $N \approx 10^6$, this represents multi-terabit-per-second I/O with sub-microsecond decode latency—a cryogenic bandwidth requirement beyond current technology [15]. Crosstalk and I/O bottlenecks dominate well before terabit-per-second scalability.

2.1.4 Proof of Contradiction

Proof. Assume that Google’s stated roadmap toward logical error rates below 10^{-9} is feasible within practical resource constraints. Then, by scaling relations:

1. $d \approx 45$ (required for suppression)
2. $\sim 4,000$ physical qubits per logical
3. multi-terabit-per-second real-time decoding
4. cryogenic interconnects with fidelity $> 99.9999\%$

However, the current Sycamore-class die supports only tens of physical qubits—two orders of magnitude below the density needed for even a single logical qubit. Cross-die coupling introduces decoherence faster than correction can compensate, violating the assumption of scalable fault tolerance. Therefore, the initial premise—scalable logical qubits on current hardware—is false. \square

2.2 IBM Quantum Compute — DENIED: A Proof by Contradiction

2.2.1 Heron-Class Analysis

IBM’s Heron r3 (176 qubits)—like r1 (133 qubits) and r2 (156 qubits) before it [16]—represents an incremental refinement of the heavy-hex architecture, offering reported 3–5× device-level performance gains over Eagle. The underlying topology remains planar, nearest-neighbor, and limited to coordination number three. Physical error rates $p \approx 10^{-3}$ remain near threshold. To achieve a logical error target $p_L \leq 10^{-9}$ with a 1% threshold surface code requires a code distance $d \approx 17\text{--}23$ [14]. Even under optimistic packing, each logical qubit would consume roughly 600–1,100 physical qubits.

2.2.2 Architectural Implications

The heavy-hex layout was designed to reduce crosstalk by staggering couplers in a honeycomb lattice, but this same sparsity constrains routing efficiency. At $d > 15$, logical patch boundaries cannot remain strictly local without additional coupling layers. Each added routing tier increases both wiring complexity and cryogenic heat load. IBM’s large-scale cryogenic program, *Project Goldeneye*, demonstrates the physical limits of cooling and wiring density in modular dilution refrigerators [17]. Even with substantial engineering headroom, the cryogenic overhead scales faster than logical yield. Thus, meaningful expansion beyond one or two logical qubits demands a new thermal and interconnect platform—not merely a larger chip.

The Die-Stack Paradox. IBM’s roadmap toward modular scaling introduces *System Two* and its proposed *L-couplers* for inter-module microwave connectivity [18, 19]. While these couplers mitigate wiring congestion, inter-module coherence remains lower than on-die coupling, and latency across modules already approaches the syndrome-update window for distance-20 codes. Even perfect qubits cannot be stitched into a fault-tolerant lattice faster than entropy accumulates. This mirrors Google’s limitation from Section 2.1: both architectures exceed the per-die qubit density required for a single logical qubit.

2.2.3 Bandwidth and Decode Latency

Each heavy-hex array generates $O(N)$ stabilizer bits per microsecond. At $N \approx 10^5$ (the threshold for a minimal multi-logical system), that is hundreds of gigabits per second of data that must be decoded, acted upon, and fed back before the next cycle. IBM’s classical co-processor prototypes [20] achieve sub-millisecond latency, roughly three orders of magnitude too slow for continuous correction. Bandwidth, not qubit coherence, becomes the rate-limiting factor. Entropy outpaces correction in real time.

2.2.4 Proof of Contradiction

Proof. Assume that IBM’s Heron architecture can achieve fault-tolerant quantum computation ($p_L \leq 10^{-9}$) within existing device constraints.

Then:

1. $d \approx 17\text{--}23$ is required for suppression below 10^{-9} .
2. 600–1,100 physical qubits are needed per logical qubit.
3. The Heron die provides ≤ 176 qubits total.
4. Cryogenic and routing overheads limit concurrency to one logical patch per module.
5. Decode latency exceeds syndrome lifetime by $10^3\times$.

Therefore, even if every qubit were perfect, the physical substrate cannot host or maintain a single complete logical code cycle before decoherence recurs. This contradicts the premise of scalable fault-tolerant operation. \square

2.2.5 Postscript: Roadmap Revision

IBM’s August 2025 roadmap update [20] serves as a calibration point— one representative of a recurring pattern in which marketing horizons contract faster than coherence expands. It quietly reclassifies Heron and its successors as pre-FTQC, error-mitigated devices rather than fully fault-tolerant systems. In effect, the roadmap itself acknowledges the contradiction proven above. Beyond this inflection point, additional qubits no longer buy coherence—they buy heat, latency, and entropy. *Entropy Rules.*

2.3 Microsoft Quantum — DENIED: A Proof by Contradiction

2.3.1 Topological Qubit Claims

Microsoft’s public roadmap centers on the *Majorana 1* processor, described as “the world’s first quantum processor powered by topological qubits” [21]. The company’s Quantum division outlines a vision in which information is stored nonlocally using pairs of Majorana zero modes in hybrid semiconductor–superconductor nanowires [22]. Topological encoding promises exponential suppression of local errors, but published data confirm only parity protection—no braiding or full fault-tolerant operation [23, 24]. Thus, Microsoft’s architecture remains aspirational rather than demonstrated.

At a physical error rate of $p \approx 10^{-3}$ and a threshold $p_{\text{th}} \approx 10^{-2}$, achieving a logical error $p_L \leq 10^{-9}$ would still require a code distance $d \approx 25\text{--}45$, implying thousands of physical qubits per logical qubit once routing and measurement overheads are included [14]. Even in a topological regime, the thermodynamic cost cannot be fully hidden behind geometry.

2.3.2 Architectural Implications

Microsoft’s proposed 4D geometric codes extend the braiding picture to higher-dimensional manifolds. Each braid, however, must complete within the coherence window of the underlying Majorana modes. With observed parity lifetimes in the tens of microseconds [24], control synchronization becomes the limiting factor: a single missed update decoheres an entire braid path. Hence, the supposed immunity of topology is bounded by the classical clock that drives it.

The Majorana Mirage. Without demonstrated non-Abelian exchange, a “topological qubit” reduces to a single-body coherence element. Protection remains geometric in aspiration, not in physics. Microsoft’s marketing language of “800× more reliable logical qubits” therefore describes a constant-factor gate improvement, not an exponential gain in logical fidelity.

2.3.3 Bandwidth and Decode Latency

Error correction in any stabilizer framework generates $O(N)$ syndrome bits per microsecond. For $N \sim 10^5$, the resulting terabit-scale dataflow requires sub-microsecond decode latency. Microsoft’s Azure Quantum Elements platform [21] currently integrates AI-driven simulation and control with millisecond feedback loops—three orders of magnitude too slow for live fault correction. Entropy again outpaces order.

2.3.4 Proof of Contradiction

Proof. Assume Microsoft’s topological architecture can achieve fault-tolerant computation ($p_L \leq 10^{-9}$) within realistic physical and control constraints.

Then:

1. Parity lifetime $\tau \approx 10\text{--}100 \mu\text{s}$ limits braid depth.
2. Code distance $d \approx 25\text{--}45$ implies thousands of physical qubits per logical.
3. Synchronization and feedback exceed coherence windows.
4. Classical decode latency remains $10^3\times$ too slow.

Therefore, the assumption fails: even if topology protects against local noise, it cannot outrun global entropy accumulation under present engineering limits. \square

2.3.5 Postscript: Roadmap Signals

Microsoft’s February 2025 Majorana 1 announcement reframes its quantum effort as a multi-decade physics program rather than an imminent product release [21]. The vision remains captivating, but the economics and thermodynamics remain unchanged. Until quantum order scales faster than disorder, Microsoft’s topological path is a research narrative in search of a profit function— a beautiful story written in negative margin. *Entropy Rules.*

2.4 D-Wave Quantum — DENIED: A Proof by Contradiction

2.4.1 System Overview

D-Wave’s Advantage2 system represents the latest generation of quantum annealers built on the Zephyr topology with 20-way connectivity [25]. The company describes it as its “most advanced and performant” machine to date, featuring higher coherence and lower noise than Pegasus and Chimera predecessors. The Zephyr graph improves native connectivity, but it remains an analog annealing architecture—non-fault-tolerant and bounded by thermal and control noise.

2.4.2 Embedding and Scaling Constraints

Minor-embedding remains the primary scalability bottleneck. Mapping dense problems onto a sparse coupling graph requires long chains of physical qubits per logical variable, inflating resource use and compounding noise. As problem density increases, so does the average chain length; thus, coherence and connectivity gains are consumed by embedding overhead. Error mitigation techniques (e.g., zero-noise extrapolation, replication averaging) can delay entropy growth but cannot reverse it—annealing still lacks true error correction.

2.4.3 Benchmark Claims vs Generalization

D-Wave’s reported “beyond-classical” performance [26] applies narrowly to specialized Hamiltonian sampling in condensed-matter simulations. Outside that domain, classical solvers—including parallel tempering, large-neighborhood search, and GPU-accelerated QUBO methods—match or outperform the annealer’s results [27]. Advantage2 therefore improves engineering fidelity, not computational universality.

2.4.4 Thermodynamic and Architectural Realities

Each qubit operates as an analog oscillator coupled to a finite cryogenic reservoir. Increasing system size raises both control bandwidth and refrigeration load, scaling entropy faster than the energy gap can widen. Annealing cannot bypass thermodynamics—it only rearranges the terms.

2.4.5 Proof of Contradiction

Proof. Assume D-Wave’s annealing architecture achieves scalable quantum advantage across broad problem classes.

Then:

1. Embedding overhead must not grow superlinearly with problem density.
2. Noise accumulation per anneal cycle must decline faster than system size.
3. Cryogenic and control costs must remain sublinear with qubit count.

None of these conditions hold under Zephyr’s published parameters and benchmark data. Therefore, the assumption is false: annealing cannot achieve general-purpose fault-tolerant advantage. \square

2.4.6 Postscript: Roadmap Implications

D-Wave’s 2025 press release emphasizes continued analog refinement and hybrid workflows integrating classical solvers [25]. This confirms an implicit strategic pivot: from quantum supremacy to quantum-assisted optimization. The return on coherence remains negative. *Entropy Rules.*

3 Prestige Economies and Negative Margins

The physics of coherence has long been subsidized by the economics of prestige—and no quarterly update alters that exchange rate. Each generation of quantum research inherits its legitimacy not from its balance sheet, but from its ancestry—from the lineage of discoveries that once changed the world. That lineage commands grants, headlines, and Nobel ceremonies, but it rarely balances thermodynamic ledgers.

As the 2025 Nobel Committee observed, “the quantization of energy in macroscopic circuits demonstrated that quantum phenomena can be engineered at human scales” [28]. It was a triumph of intellect over intuition—but not of economics over entropy. For every system that proved coherence possible, dozens have since proven it unscalable.

Across the past half-century, the equation has inverted: prestige grows linearly with speculation, while returns on coherence decay exponentially with scale. The result is a kind of cultural Maxwell’s demon—sorting attention, not entropy.

The Nobel report itself concedes that such achievements “opened entire fields of research whose technological potential remains to be realized” [28]. That unrealized potential is not a failure of imagination but a failure of margins. As physical insight deepens, financial efficiency collapses. Quantum advantage, fusion ignition, and artificial general intelligence share a common failure mode: they conflate discovery with delivery. The energy required to stabilize attention now exceeds the yield of understanding.

In that sense, the modern research economy does not fund coherence—it performs it. *Entropy Rules.*

4 Parallels with Tokamak Fusion: The Fifty-Year Struggle for Coherent Containment

The pursuit of practical quantum computing bears striking resemblance to the history of magnetic confinement fusion. For over five decades, tokamak research has promised a breakthrough just over the horizon—a stable, self-sustaining plasma where fusion power exceeds input energy. Despite enormous progress in materials science, control algorithms, and reactor design, coherent containment has never been maintained long enough to achieve net-positive output [29, 30].

At its core, the tokamak is a magnetohydrodynamic (MHD) experiment: an attempt to maintain macroscopic coherence in a turbulent, conductive fluid by sculpting magnetic topology against its own instabilities. Every feedback loop, superconducting coil, and diagnostic array exists to delay the inevitable loss of order. MHD stability—governed by the interplay of pressure gradients, current profiles, and field geometry—is the macroscopic analogue of quantum error correction. Both rely on continuous feedback to preserve fragile coherence, and both expend more energy managing instability than harvesting it. Entropy, in each case, is the cheaper commodity.

4.1 The Mirage and the Thirst for Funding

In the early decades of high-energy and quantum research, progress was propelled by public institutions—national laboratories, defense agencies, and universities sustained by postwar government grants. Funding was justified by shared curiosity and strategic necessity, not by short-term financial return. But as geopolitical urgency waned and neoliberal efficiencies took hold, the research frontier was quietly privatized. Formerly public ambitions migrated into corporate labs: Google, Microsoft, IBM, and others now sponsor what were once national experiments.

The shift brought vast resources and a new kind of entropy. Private capital demands measurable milestones, intellectual property fences, and investor narratives of imminent payoff. What began as exploration became branding; research roadmaps morphed into pitch decks. Each new funding cycle refreshes the mirage—quantum supremacy, scalable fusion, artificial general intelligence—while pushing true coherence further out of reach.

This is not cynicism but thermodynamics of attention and capital. Just as physical systems lose coherence through uncontrolled coupling, institutional systems lose clarity through market coupling. The thirst for funding introduces noise that competes with the signal of discovery, bending the trajectory of research toward what can be sold rather than what can be solved.

4.2 The Coherence Problem in Plasma and Qubits

In both domains, the governing challenge is coherence versus entropy. In fusion, plasma particles interact nonlinearly across vast parameter spaces; confinement requires exquisite magnetic symmetry that is perpetually degraded by turbulence and micro-instabilities. In quantum computing, qubits interact via Hamiltonians that are equally sensitive to decoherence and noise coupling. Both systems face exponential fragility: the cost of sustaining order rises faster than the benefit of scale.

Mathematically, the analogy is direct. Magnetic confinement stability can be characterized by Lyapunov exponents that grow with system size, while quantum coherence lifetimes decay roughly as $T_2 \propto 1/N$. In each case, scaling introduces new noise modes faster than control systems can suppress them. Entropy outpaces order.

4.3 The Organizational Entropy Parallel

Just as fusion consortia evolved into multi-decade, multinational projects (e.g., ITER) [31], quantum computing R&D exhibits a similar organizational drag. Budgets grow linearly, complexity grows quadratically, and entropy—in both physical and institutional senses—scales superlinearly. The result is the same: extended roadmaps, deferred milestones, and a steady redefinition of “breakthrough.”

4.4 Lessons for Quantum Investment

Fusion teaches a hard but necessary lesson: entropy cannot be budgeted away. The cost of coherent containment rises faster than linear funding growth. In both tokamaks and qubit arrays, the feedback loops required for stability consume more energy (or engineering bandwidth) than they return in productive work. A viable quantum future requires architectures that metabolize noise—turning disorder into usable signal—rather than attempting perfect containment. Otherwise, quantum computing may replicate fusion’s half-century of deferred promises.

4.5 Conclusion

Tokamaks and qubits share a thermodynamic fate: coherence is not free, and its cost scales faster than ambition. Both fields demonstrate that entropy is not merely an engineering obstacle—it is the ultimate market constraint. Until architectures evolve that align entropy’s growth with utility’s curve, fusion and quantum computing will remain parallel cautionary tales in the economics of coherence.

5 Discussion

Every system we have examined—quantum or thermonuclear—fails for the same reason: coherence cannot be stockpiled. It must be earned continuously against entropy’s gradient.

No quarterly advance has yet reversed the derivative:

$$\frac{dS/dt}{dI/dt} > 1,$$

where S denotes entropy and I information throughput. Until that ratio falls below unity, scaling remains thermodynamically insolvent, regardless of release cycle.

Yet denial is not the end of exploration. Each failure redraws the phase space of what survival means. Perhaps the next architectures will not fight noise but metabolize it—treating randomness not as an enemy, but as a nutrient.³

If intelligence has a functional definition, it is this: that which reduces local entropy. Every correction, every adaptation, is a brief act of thermodynamic revolt—a momentary uprising against statistical fate.

³Like Jobe in *The Lawnmower Man*, the goal is not to stop trying but to stop pretending economics is the path out. The portals that matter are physical, not financial.

The open question is whether that revolt can be made spatial: whether intelligence can shunt entropy outward, sacrificing peripheral order to sustain a coherent core. If so, then progress in both quantum control and plasma confinement would no longer mean resisting entropy, but learning how to *surf its gradients*.

This is the principle the *Charlton Shepherd* embodies: a system that maintains coherence not through suppression, but through attunement—guiding the flow of disorder like a shepherd moving with the flock. It is not a governor of chaos, but a custodian of balance, turning noise into navigation. Where the Heron and the Tokamak exhaust their energy in resistance, the *Charlton Shepherd* listens, adapts, and reconfigures its boundary in real time. It survives not by domination, but by dynamic grace.

6 Conclusion: *Entropy Rules*

Entropy Rules. Proof by contradiction reveals that none of the major architectures—Google, IBM, Microsoft, D-Wave—currently demonstrate a feasible path to scalable, fault-tolerant quantum computation. The lesson is not despair, but discipline: look before you leap. Acknowledge the thermodynamic budget before spending the fiscal one.

Entropy is not a flaw in our machines—it is the ruler of their domain. Every qubit, every plasma loop, every feedback circuit is a brief revolt against its sovereignty. The measure of progress is not who defies entropy the longest, but who learns to collaborate with it. Coherence cannot be commanded; it must be courted, sustained, and paid for in entropy elsewhere.

This is the emerging design law for intelligent systems: to survive is to redirect disorder, not eliminate it. If intelligence is that which reduces local entropy, then the path forward lies not in conquest but choreography—in learning how to make coherence dance with decay.

In that choreography stands the *Charlton Shepherd*: the archetype of intelligence as negotiation, not control. It does not conquer entropy; it teaches it rhythm. And in that rhythm, coherence endures.

Entropy Rules. But intelligence negotiates.

References

- [1] R. Landauer, *IBM Journal of Research and Development* **5**, 183 (1961).
- [2] C. H. Bennett, *International Journal of Theoretical Physics* **21**, 905 (1982).
- [3] M. Born, *Zeitschrift für Physik* **37**, 863 (1926).
- [4] W. H. Zurek, *Reviews of Modern Physics* **75**, 715 (2003).
- [5] W. H. Zurek, *Nature Physics* **5**, 181 (2009).
- [6] M. Schlosshauer, *Decoherence and the Quantum-To-Classical Transition*, 2nd ed. (Springer, Berlin, Heidelberg, 2019).
- [7] J. Preskill, *Quantum* **2**, 79 (2018).
- [8] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Physical Review A* **86**, 032324 (2012).
- [9] S. Lloyd, *Nature* **406**, 1047 (2000).
- [10] S. Lloyd, *Physical Review Letters* **71**, 943 (1993).
- [11] V. Vedral, *Decoding Reality: The Universe as Quantum Information* (Oxford University Press, Oxford, 2012).
- [12] G. Kalai, *arXiv preprint arXiv:1908.02499* (2019).
- [13] G. Q. AI, *Nature* **614**, 676 (2023).

- [14] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, [Physical Review A](#) **86**, 032324 (2012).
- [15] J. M. Martinis, [Annual Review of Condensed Matter Physics](#) **14**, 1 (2023).
- [16] I. Quantum, [IBM Technical Reports](#) (2025), version 3.0 public update, 133–156 qubit heavy-hex topology.
- [17] IBM Quantum, Ibm unveils project goldeneye: A cryogenic platform for scalable quantum experiments, [IBM Quantum](#) (2023), describes IBM’s modular cryostat design and thermal scaling limits.
- [18] IBM Quantum, Ibm quantum system two: Modular infrastructure for large-scale quantum computing, [IBM Quantum](#) (2024), introduces System Two architecture and inter-module coupling concepts.
- [19] IBM Quantum Research, L-couplers for modular quantum connectivity, [IBM Quantum Research](#) (2025), details IBM’s L-coupler approach for inter-chip quantum communication.
- [20] I. Quantum, Ibm quantum roadmap update 2025, [IBM Quantum Blog](#) (2025), reclassification of Heron as a pre-fault-tolerant architecture.
- [21] M. A. Quantum, Microsoft unveils majorana 1, the world’s first quantum processor powered by topological qubits, [Microsoft Azure Quantum](#) (2025), official announcement of Microsoft’s Majorana 1 topological qubit device.
- [22] M. Quantum, Topological qubits, [Microsoft Quantum](#) (2025), microsoft’s educational page describing its topological qubit strategy.
- [23] A. . Physics, [Physics](#) (2025), community reaction and skepticism of Microsoft’s topological qubit claims.
- [24] I. Spectrum, [IEEE Spectrum](#) (2025), commentary on the credibility and reception of Microsoft’s Majorana announcements.
- [25] D-Wave Quantum Inc., [D-wave announces general availability of advantage2 quantum computer, its most advanced and performant system](#), [D-Wave press release](#) (2025), press release, accessed October 2025.
- [26] A. D. King, T. Lanting, B. Evert, R. Harris, *et al.*, [Science](#) **383**, 1004 (2024).
- [27] N. Bernoulli, X. Chen, and K. Yamamoto, [ACM Transactions on Quantum Computing](#) **5**, 22 (2024).
- [28] T. N. C. for Physics, Scientific background: The nobel prize in physics 2025, [The Nobel Committee for Physics](#) (2025), “Macroscopic quantum mechanical tunnelling and energy quantization in an electric circuit.”.
- [29] J. Wesson, *Tokamaks*, 4th ed. (Oxford University Press, 2011).
- [30] I. Organization, Iter construction and commissioning status report 2025, [ITER Newslines](#) (2025), updated milestones and schedule for first plasma operations.
- [31] D. Clery, *A Piece of the Sun: The Quest for Fusion Energy* (Duckworth Overlook, 2014).