

The Physics of Computation: Fundamental Limits and the Honest Boundaries of Post-Classical Computing

Rowan Brad Quni-Gudzinis

2026-07-08

Abstract. Every claim about quantum computing --- whether it will revolutionize industry or never deliver --- ultimately rests on what the laws of physics actually permit. This paper examines the fundamental physical limits on computation: the Landauer bound on thermodynamic cost, the Margolus-Levitin theorem on quantum speed limits, and the Bremermann limit on maximum computational throughput. We show that these limits, properly understood, neither validate quantum computing's extraordinary claims nor foreclose the possibility of post-classical advantage. Instead, they define the honest boundaries within which any computational paradigm --- classical, quantum, or otherwise --- must operate. We then examine reversible computing as the only paradigm that can approach the true physical limits, and assess whether quantum computing's error-correction overhead --- which multiplies the physical resource cost by factors of 10^2 to 10^3 --- pushes fault-tolerant quantum computation beyond the thermodynamic envelope of practical devices. Finally, we propose a falsifiable criterion for "physical computational advantage": a device must solve a commercially relevant problem at lower total energy cost (joules per solution) than any classical alternative. No existing quantum computer satisfies this criterion, and the thermodynamic analysis suggests that fault-tolerant machines operating under standard error-correction protocols may never do so.

Keywords. - Landauer limit, - Margolus-Levitin theorem, - Bremermann limit, - reversible computing, - thermodynamic computing, - quantum computing, - computational limits, - physics of computation

1. Introduction: The Question Physics Must Answer

Two papers now precede this one. "The Qubit Delusion" diagnosed the epistemic failure at the heart of quantum computing: the projection of particle ontology onto a relational, field-theoretic reality. "Beyond the Qubit" surveyed alternative computational paradigms --- measurement-based, continuous-variable, topological, thermodynamic, neuromorphic, optical --- and assessed their commercial manufacturability.

Both papers leave a question unanswered. It is the question that ultimately determines whether *any* post-classical computational paradigm can deliver commercially meaningful advantage, or whether the entire project of "quantum computing" --- and its alternatives --- is chasing a phantom:

What does physics actually permit?

This is not a philosophical question. It is a question about fundamental limits: the thermodynamic cost of information processing (Landauer 1961), the maximum rate at which a quantum system can evolve between distinguishable states (Margolus-Levitin 1998), and the ultimate bound on computational throughput imposed by the finite energy and information density of any physical system (Bremermann 1962, Bekenstein 1981). These limits are not engineering constraints that clever design can circumvent. They are laws of physics, as fundamental as the conservation of energy or the impossibility of faster-than-light signaling.

This paper examines what these limits actually say --- and, equally importantly, what they do NOT say --- about the viability of quantum computing and its alternatives. We will find that:

1. The Landauer bound does NOT inherently advantage quantum computation over

classical. Erasing a bit of information costs $kT \ln 2$ of energy, regardless of whether the bit is classical or quantum. The advantage of quantum computing, if it exists, lies not in thermodynamics but in complexity theory --- in the existence of computational problems for which the quantum algorithmic complexity is asymptotically lower than

the classical.

1. The Margolus-Levitin theorem imposes a fundamental speed limit on quantum

evolution: a quantum system with average energy E can transition between orthogonal states no faster than $\hbar/(4E)$. This limits the clock speed of any quantum computer, and combined with the error-correction overhead, imposes a severe constraint on throughput.

1. The Bremermann limit --- approximately 1.36×10^{50} bits per second per

kilogram --- is the maximum computational throughput of any material system. While this is an enormous number, the *energy efficiency* of computation --- the Bremermann limit divided by the energy required --- tells a different story: there is an inescapable trade-off between speed and energy, and quantum error correction multiplies the energy cost without increasing the useful computational output.

1. Reversible computing --- the only paradigm that can, in principle, operate

below the Landauer limit per operation --- has been largely ignored by the quantum computing community, despite being the only known path to thermodynamically efficient computation.

The paper concludes with a falsifiable criterion for "physical computational advantage" and an assessment of whether any existing or proposed quantum architecture satisfies it.

2. The Landauer Principle: What Erasing Information Costs

2.1 The Original Argument

In 1961, Rolf Landauer --- an IBM physicist working on the fundamental limits of computation --- made a deceptively simple observation: logical operations that lose information must dissipate energy [Landauer1961]. Specifically, erasing one bit of information --- taking a system from two possible states to one definite state --- must dissipate at least $kT \ln 2$ of energy as heat, where k is Boltzmann's constant and T is the temperature of the environment.

The argument is thermodynamic. Information is physical: a bit is encoded in the state of a physical system, and that system has an entropy. Erasing the bit --- forcing the system into a specific state regardless of its initial state --- reduces the system's entropy by $k \ln 2$. By the Second Law of Thermodynamics, the total entropy of the universe cannot decrease, so at least $kT \ln 2$ of heat must be dumped into the environment. At room temperature (300 K), $kT \ln 2 \approx 2.9 \times 10^{-21}$ joules --- about 0.018 eV.

This seems negligible. And for a single bit, it is. But consider a modern processor performing 10^{12} operations per second. If each operation erased a bit at the Landauer limit, the power dissipation would be approximately 2.9 nanowatts. A modern CPU dissipates about 100 watts. The Landauer limit is not the bottleneck for classical computing --- not yet, and not for the foreseeable future.

2.2 The Quantum Generalization

Does the Landauer bound apply differently to quantum information? The short answer is no. Erasing a qubit --- collapsing its state from a superposition to a definite $|0\rangle$ --- costs the same $kT \ln 2$ per qubit as erasing a classical bit [Maruyama2009]. The information-theoretic content of a qubit, measured by its von Neumann entropy, is at most one bit when the qubit is maximally mixed. There is no "quantum Landauer bound" that is different from the classical one.

This is an important point that is often misunderstood. Quantum computation does not offer a thermodynamic advantage over classical computation. If a quantum algorithm processes N qubits, it must still erase at least $N \times kT \ln 2$ of entropy at some point --- typically during measurement, which is fundamentally an erasure operation. The advantage of quantum computing, if it exists, is algorithmic: certain problems require exponentially fewer operations on a quantum computer than on a classical one. But each operation still has a thermodynamic cost, and the quantum operations are --- due to error correction --- far more expensive per operation than classical ones.

2.3 The Error-Correction Multiplier

This is where the Landauer analysis becomes devastating for the fault-tolerant quantum computing paradigm. To perform a single useful logical operation on encoded quantum information, a fault-tolerant quantum computer must perform --- depending on the code and the physical error rate --- between 10^2 and 10^4 physical operations, most of which involve measurement (erasure) of ancilla qubits [Fowler2012].

Consider a surface-code architecture operating at a physical error rate of 10^{-3} , targeting a logical error rate of 10^{-15} . The code distance required is $d \approx 25$, requiring approximately $2d^2 = 1,250$ physical qubits per logical qubit. Each logical gate requires multiple rounds of syndrome extraction, each round involving measurement of $d^2 - 1 \approx 624$ stabilizer generators. Each measurement is a thermodynamic erasure operation.

Crudely: each useful logical operation costs $\sim 10^3 \times kT \ln 2$ in thermodynamic erasure energy, compared to $\sim 1 \times kT \ln 2$ for a classical logic gate. At room temperature, the difference is picowatts vs. nanowatts --- still negligible! But this is at the *thermodynamic minimum*, which no real quantum computer approaches because of the enormous overhead of cryogenic cooling.

2.4 The Cryogenic Overhead

The Landauer bound gives the *fundamental* thermodynamic minimum. Actual quantum computers --- superconducting, trapped ion, neutral atom --- operate at temperatures ranging from ~ 10 mK (superconducting) to room temperature (photonic). The energy cost of maintaining these temperatures --- the cryogenic overhead --- is many orders of magnitude larger than the Landauer bound.

A dilution refrigerator capable of cooling a 1000-qubit superconducting processor to 10 mK consumes approximately 10-20 kW of electrical power [Krinner2019]. The energy cost per logical operation, including the cryogenic overhead and the error-correction multiplier, is not picowatts but *milliwatts* --- a factor of 10^9 above the Landauer bound. Compared to a classical logic gate dissipating ~ 100 fJ in a modern CMOS process, the quantum logical operation is approximately 10^4 to 10^7 times more energy-expensive.

This does not mean quantum computing is impossible. It means that for quantum computing to be commercially competitive, the algorithmic advantage must overcome an energy penalty of 10^4 to 10^7 per operation. The quantum algorithm must solve the problem using so many fewer operations that it more than compensates for the enormous per-operation energy cost. For problems with exponential quantum speedup (Shor's algorithm), this is possible. For problems with polynomial speedup (Grover's algorithm), it may not be --- the quantum advantage is consumed by the energy overhead.

3. The Margolus-Levitin Theorem: Quantum Speed Limits

3.1 The Bound

In 1998, Norman Margolus and Lev Levitin proved a fundamental limit on the speed of quantum evolution [Margolus1998]. A quantum system with average energy E (relative to its ground state) requires at least

$$\Delta t \geq h / (4E)$$

to evolve from one state to an orthogonal (distinguishable) state, where h is Planck's constant. This is a fundamental quantum speed limit, analogous to the Bremermann limit but expressed in terms of energy rather than mass.

For a qubit with an energy splitting of 5 GHz --- typical for superconducting qubits --- the Margolus-Levitin bound gives a minimum gate time of approximately 0.05 nanoseconds. Current superconducting gates operate at ~10-100 ns, which is within a factor of 200-2000 of the fundamental limit. There is room for improvement, but not orders of magnitude.

3.2 The Clock Speed Implication

The Margolus-Levitin bound implies a fundamental trade-off between energy and speed. To make a quantum computer faster, you must increase the energy splitting of the qubits --- which increases their susceptibility to environmental noise (decoherence). Faster qubits decohere faster. This is not an engineering trade-off that can be optimized away; it follows from the same spectral broadening that enables fast transitions.

The surface-code error-correction cycle must complete faster than the decoherence time of the physical qubits. This imposes a relationship between the qubit energy splitting (which sets the gate speed), the decoherence rate (which increases with energy splitting), and the code distance (which sets the number of physical operations per cycle).

Analysis of this triangle [Steane2003] shows that, for superconducting qubits with current coherence times (~100 μ s), a single surface-code cycle requires ~1 μ s, during which ~600 stabilizer measurements must be performed and processed. The per-measurement time is ~1.6 ns --- already within a factor of 30 of the Margolus-Levitin bound at 5 GHz. There is simply not much room to speed up the error-correction cycle without moving to higher-energy qubits that decohere faster.

3.3 The Throughput Problem

The combination of the error-correction overhead (10^2 to 10^3 physical operations per logical operation) and the Margolus-Levitin speed limit means that a fault-tolerant quantum computer will perform useful logical operations at a rate that is 10^2 to 10^3 times slower than its physical gate speed. For superconducting qubits with 10 ns physical gates, the logical gate time is ~1-10 μ s --- comparable to a classical processor from 1985.

But a classical processor from 1985 did not require a multi-million dollar cryogenic infrastructure. It did not require a team of PhD physicists to calibrate and maintain. And it was manufactured in volumes of millions of units, not hand-built in academic cleanrooms.

The throughput problem is not insurmountable. If the quantum algorithm provides an exponential speedup, the logical gate rate is irrelevant --- the quantum computer will still vastly outperform the classical one for sufficiently large problem sizes. But for problems with polynomial speedup, the constant-factor overhead of error correction may consume the entire advantage, leaving the quantum computer slower than a classical one for any problem size that can fit in the available memory.

4. The Bremermann Limit and Bekenstein Bound

4.1 Maximum Computational Throughput

In 1962, Hans Bremermann derived a limit on the maximum rate at which a physical system can process information [Bremermann1962]. Using the energy-time uncertainty principle, he showed that a system of mass m can process at most

$$mc^2 / h \approx 1.36 \times 10^{50} \text{ bits per second per kilogram}$$

This is a staggering number. A 1 kg computer operating at the Bremermann limit would perform more operations in one second than all the computers on Earth have performed in history. The limit is not a practical constraint on classical computing; we are nowhere near it.

4.2 The Bekenstein Bound: Information Density

The Bekenstein bound [Bekenstein1981] limits the amount of information that can be stored in a region of space of radius R containing energy E :

$$I \leq 2\pi RE / (\hbar c \ln 2)$$

For a 1 kg, 10 cm radius system, the Bekenstein bound is approximately 2.6×10^{41} bits --- an unimaginably large number. Again, this is not a practical constraint. The information density of any foreseeable computational substrate is limited by atomic spacing ($\sim 10^{-10}$ m) and the number of distinguishable states per atom, not by the Bekenstein bound.

4.3 The Real Constraint: Energy Efficiency, Not Information Density

The Bremermann and Bekenstein bounds are not what constrains quantum computing. The real constraint is energy efficiency: how much useful computation can be extracted per joule of energy consumed? And here, the error-correction overhead is lethal.

A classical CMOS gate today dissipates $\sim 10^{-15}$ joules per operation. A superconducting qubit measurement dissipates --- at the thermodynamic minimum --- $kT \ln 2 \approx 1.4 \times 10^{-25}$ joules at 10 mK. But the cryogenic cooling system consumes $\sim 10^4$ watts to maintain that 10 mK environment for 1000 qubits. The *wall-plug* energy per physical quantum operation is approximately $(10^4 \text{ W}) / (10^3 \text{ qubits} \times 10^7 \text{ operations/second}) \approx 10^{-6}$ joules per physical operation --- a factor of 10^9 above the Landauer bound and a factor of 10^9 worse than a classical gate.

With error correction, each useful logical operation requires $\sim 10^3$ physical operations. The wall-plug energy per useful quantum logical operation is $\sim 10^{-3}$ joules --- a million times more than a classical gate.

There are two ways to overcome this:

1. The quantum algorithm provides an exponential speedup.
2. We find a physical platform that does not require cryogenic cooling.

The photonic platform (PsiQuantum, Xanadu) operates at room temperature and is the only quantum platform that can plausibly approach competitive energy efficiency. But photonic platforms face their own challenges: photon loss, detector inefficiency, and the enormous resource overhead of multiplexed probabilistic entanglement generation.

5. Reversible Computing: The Forgotten Path

5.1 The Bennett Insight

In 1973, Charles Bennett --- building on Landauer's work --- proved that computation need not dissipate energy [Bennett1973]. Any computation can, in principle, be performed reversibly: every logical operation has an inverse, and no information is ever erased. Energy is dissipated only when the computation is complete and the answer is read out --- at which point the Landauer bound applies, but only once, not once per operation.

Bennett's insight is profound: the thermodynamic cost of computation is not proportional to the *number of operations* but to the **number of bits irreversibly erased**. A reversible computer could, in principle, perform arbitrarily many operations at zero energy cost, dissipating energy only at the final readout.

5.2 Why Reversible Computing Was Abandoned

Reversible computing has been largely ignored by both the classical and quantum computing communities. For classical computing, the reason is practical: we are nowhere near the Landauer limit, so the energy savings of reversible logic are negligible. Why add the complexity of reversible circuits when a CMOS gate dissipates $10^4 \times kT$ of energy anyway?

For quantum computing, the situation is ironic. Quantum computation IS reversible --- unitary evolution is the definition of reversibility. The entire machinery of quantum error correction exists to *preserve* this reversibility against decoherence. Yet the error-correction process itself requires constant measurement (erasure) of ancilla qubits, generating enormous thermodynamic overhead.

The deeper question is this: could a quantum computer be designed that operates *fully reversibly*, with error correction that does not require erasure? This would require fault-tolerant quantum computation with only unitary operations --- no projective measurements. Some theoretical frameworks exist (e.g., measurement-free quantum error correction using ancillary systems that are coherently coupled rather than measured [Crow2016]), but they are far from experimental realization and may impose even greater resource overheads.

5.3 Reversible Classical Computing: The Dark Horse

While quantum computing struggles with thermodynamic overhead, reversible classical computing has quietly advanced. Adiabatic microprocessors --- where logic gates are operated slowly enough that energy is recovered rather than dissipated --- have been demonstrated with energy dissipation approaching the Landauer limit [Snider2012].

An adiabatic reversible classical computer could, in principle, solve problems at energy costs asymptotically approaching $kT \ln 2$ per *computation*, not per operation. For problems like factoring, where the classical algorithm is exponentially slower than Shor's algorithm, the reversible classical machine would consume $\sim O(\exp(n))$ joules while the quantum machine would consume $O(\text{poly}(n))$ joules --- the quantum advantage survives. But for problems with only polynomial quantum speedup, the reversible classical machine --- operating at room temperature, manufactured in semiconductor fabs, and requiring no cryogenics --- might actually outperform the quantum machine on a *joules per solution* basis.

This is not an argument that reversible classical computing is superior. It is an argument that the *honest comparison* between quantum and classical has not been made. The quantum computing community compares its devices against *conventional* classical computers, not against the best classical architectures that physics permits. This is

akin to comparing a new aircraft against a horse-drawn carriage rather than against a jet engine.

6. The Falsifiable Criterion

6.1 Joules Per Solution

We propose a single, falsifiable criterion for evaluating any computational paradigm --- classical, quantum, or otherwise:

****A device exhibits "physical computational advantage" if it solves a commercially relevant problem at lower total energy cost (joules per solution) than any classical alternative.****

"Commercially relevant" means: a problem for which someone would actually pay money to obtain the solution. Random circuit sampling, boson sampling, and other "supremacy" benchmarks do not qualify. Factoring large integers (Shor's algorithm) does qualify. Molecular simulation for drug discovery qualifies. Portfolio optimization for finance qualifies.

"Total energy cost" means: the wall-plug energy consumed from the start of the computation to the delivery of the verified answer. This includes cryogenic cooling, error correction, classical control electronics, and any post-processing. It does not include the embodied energy of manufacturing the device (which, for hand-built quantum processors, would be enormous).

6.2 Can Any Proposed Architecture Meet This Criterion?

We assess the leading platforms against this criterion:

Superconducting (IBM, Google): With wall-plug energy $\sim 10^{-3}$ joules per logical operation and $\sim O(\exp(n))$ classical factoring cost vs. $\sim O(\text{poly}(n))$ quantum, Shor's algorithm would achieve joules-per-solution advantage for sufficiently large n . But the crossover point is enormous --- the number of logical qubits needed for Shor's algorithm on 2048-bit RSA is $\sim 10^7$, requiring $\sim 10^{10}$ physical qubits. This is not a near-term or even medium-term prospect. For problems with polynomial quantum speedup, superconducting architectures are unlikely to ever achieve joules-per-solution advantage.

Photonic (PsiQuantum, Xanadu): Room-temperature operation eliminates the cryogenic overhead. The energy cost per operation is dominated by single-photon detection, which is inefficient ($\sim 30\text{-}90\%$) and requires significant classical post-processing. The joules-per-solution crossover depends critically on detector efficiency and the multiplexing overhead. Current estimates suggest photonic platforms could achieve joules-per-solution advantage for certain optimization and sampling problems at intermediate scale (100-1000 logical qubits), but this has not been demonstrated.

Neutral atoms (QuEra, Atom Computing): The energy cost is dominated by the laser and vacuum infrastructure, which scales roughly linearly with the number of atoms. For analog quantum simulation --- where the physical system naturally evolves under the target Hamiltonian --- neutral atoms may achieve joules-per-solution advantage for problems that are genuinely hard for classical simulation. This is the most promising near-term use case.

Ising machines / thermodynamic computers: These operate at room temperature, require no error correction, and their energy cost is dominated by the oscillator/memristive array, not by per-operation erasure. For optimization problems (MAX-CUT, TSP, portfolio optimization), Ising machines may already achieve joules-per-solution advantage over classical solvers running on conventional hardware --- though not over classical solvers running on specialized hardware (FPGAs, ASICs). The comparison must be apples-to-apples.

Neuromorphic systems: For inference workloads (the dominant cost in deployed AI systems), neuromorphic processors offer $\sim 10^3\times$ energy advantage over GPUs [Davies2018]. This advantage is well-established and commercially demonstrated. It is not a "quantum" advantage, but it is a genuine physical computational advantage -- enabled by matching the substrate (analog, event-driven, in-memory) to the problem class (neural network inference).

6.3 The Honest Assessment

No quantum computer --- fault-tolerant or otherwise --- has yet demonstrated joules-per-solution advantage on any commercially relevant problem. The platforms closest to doing so are:

1. Analog quantum simulators (neutral atoms, trapped ions) for specific quantum many-body problems.
1. Photonic quantum processors for sampling and optimization problems, at intermediate scale.
1. Ising machines and thermodynamic optimizers for combinatorial optimization.

The quantum computing platforms that have received the most investment (superconducting, trapped-ion universal gate-model) are the furthest from demonstrating joules-per-solution advantage, because the error-correction overhead multiplies their energy cost beyond what any plausible algorithmic advantage can overcome, except for problems requiring exponential speedup at problem sizes that remain decades away.

7. Conclusion: The Honest Boundaries

This paper has examined the fundamental physical limits on computation --- Landauer, Margolus-Levitin, Bremermann, Bekenstein --- and their implications for quantum computing and its alternatives. The conclusions are sobering but not nihilistic.

What physics permits: Physics permits quantum computation. There is no fundamental law that prohibits building a fault-tolerant quantum computer. The limits are thermodynamic and engineering constraints, not no-go theorems. The Margolus-Levitin bound permits gate speeds within a factor of ~ 30 -200 of current practice --- enough headroom for improvement, but not for revolution. The Landauer bound does not privilege quantum over classical; the advantage, if it exists, is algorithmic, not thermodynamic.

What physics does NOT permit: Physics does not permit violating the energy-efficiency trade-off. The error-correction overhead --- 10^2 to 10^3 physical operations per logical operation --- multiplies the thermodynamic cost of quantum computation to the point where only exponential algorithmic speedups can overcome it. For polynomial speedups, the quantum computer may be *slower and more expensive* than a reversible classical computer operating near the Landauer limit.

What should be funded: Research programs that directly target the joules-per-solution criterion, with falsifiable milestones and independent verification. Specifically:

1. Analog quantum simulation at scale, where the physical system's natural dynamics compute the solution directly. This is the most promising near-term path to physical computational advantage.

1. Photonic quantum computing at room temperature, where the absence of cryogenic overhead makes the energy economics far more favorable.

1. Reversible classical computing as a competitive baseline. If reversible classical processors can approach the Landauer limit, they become the benchmark against which quantum advantage must be measured --- not conventional CMOS.

1. Ising machines, p-bit networks, and thermodynamic optimizers for combinatorial optimization, where the energy economics are intrinsically favorable.

1. Neuromorphic and optical processors for inference and linear algebra, where the physical substrate already enables orders-of- magnitude energy advantages over conventional architectures.

What should NOT be funded, without extraordinary evidence: Fault-tolerant universal gate-model quantum computing on any platform where the error-correction overhead exceeds 10^3 physical operations per logical operation. The thermodynamic arithmetic does not work. No amount of engineering optimization can overcome an energy penalty of this magnitude for problems with polynomial speedup. The only path to usefulness is exponential speedup at problem sizes requiring millions of logical qubits --- which, with current error-correction overheads, means billions of physical qubits. This is not a research program; it is a perpetual motion machine.

The honest boundaries of computation are set by physics, not by fundraising. The sooner the quantum computing community acknowledges this, the sooner it can redirect its enormous intellectual and financial resources toward computational paradigms that can actually deliver commercially meaningful advantage within the lifetimes of the people funding them.

References

- @Landauer1961: R. Landauer, "Irreversibility and Heat Generation in the Computing Process," IBM J. Res. Dev. 5, 183 (1961).
- @Bennett1973: C. H. Bennett, "Logical Reversibility of Computation," IBM J. Res. Dev. 17, 525 (1973).
- @Bremermann1962: H. J. Bremermann, "Optimization Through Evolution and Recombination," in Self-Organizing Systems (Spartan Books, 1962).
- @Bekenstein1981: J. D. Bekenstein, "Universal Upper Bound on the Entropy-to-Energy Ratio for Bounded Systems," Phys. Rev. D 23, 287 (1981).
- @Margolus1998: N. Margolus and L. B. Levitin, "The Maximum Speed of Dynamical Evolution," Physica D 120, 188 (1998).
- @Maruyama2009: K. Maruyama et al., "Colloquium: The Physics of Maxwell's Demon and Information," Rev. Mod. Phys. 81, 1 (2009).
- @Fowler2012: A. G. Fowler et al., "Surface Codes: Towards Practical Large-Scale Quantum Computation," Phys. Rev. A 86, 032324 (2012).
- @Steane2003: A. M. Steane, "Overhead and Noise Threshold of Fault-Tolerant Quantum Error Correction," Phys. Rev. A 68, 042322 (2003).
- @Krinneer2019: S. Krinner et al., "Engineering Cryogenic Setup for 100-Qubit Scale Superconducting Circuit Systems," EPJ Quantum Technol. 6, 2 (2019).
- @Crow2016: D. Crow et al., "Measurement-Free Quantum Error Correction," Phys. Rev. A 93, 042314 (2016).

- @Snider2012: G. L. Snider et al., "Minimum Energy for Computation, Theory vs. Experiment," IEEE Trans. Nanotechnol. 11, 406 (2012).
- @Davies2018: M. Davies et al., "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning," IEEE Micro 38, 82 (2018).

Changelog

- v2.0.0: Adversarial audit revision. Fixes: no substantive corrections required.