

# Beyond the Qubit: Constructive Paradigms for Post-Particle Computation

Rowan Brad Quni-Gudzinis

2026-07-08

**Abstract.** If "The Qubit Delusion" diagnosed the epistemic failure at the heart of quantum computing --- the projection of particle ontology onto a relational, field-theoretic reality --- then this companion paper asks the constructive question: what comes next? We begin with an ontological stripping: what does quantum mechanics actually demand of a computational substrate, once all anthropocentric and classical baggage is removed? From this foundation we survey alternative computational paradigms (measurement-based, continuous-variable, topological, field-theoretic, analog) that are more faithful to quantum reality than the qubit-gate-circuit scaffold. We then assess which of these are commercially manufacturable --- not in theory, but under the actual constraints of materials, thermodynamics, and industrial economics. Finally, we ask whether "quantum computing" itself is the right frame at all, or whether the deeper lesson --- that computation is a physical process and the substrate IS the algorithm --- points toward wholly different computational architectures: thermodynamic, neuromorphic, optical, and biological systems where the physics does the computing directly rather than being force-fit into a discrete symbolic scaffold.

**Keywords.** - quantum computing, - alternative computing architectures, - relational quantum mechanics, - measurement-based quantum computing, - neuromorphic computing, - thermodynamic computing, - quantum foundations, - philosophy of computation

## 1. The Question Behind the Critique

"The Qubit Delusion" advanced a diagnosis: the quantum computing industry's failure to deliver commercially viable machines is not primarily an engineering problem but an epistemic one. The qubit-gate-circuit model imports a particle ontology into quantum mechanics --- as if qubits were little quantum billiard balls that we poke with gates --- and this ontology is inconsistent with quantum field theory, relational quantum mechanics, and the physics of continuous, correlated quantum systems.

But diagnosis without prescription is incomplete. If the qubit is a scaffold, not an invariant, then what IS the invariant? If the gate-circuit model is a map-territory confusion, what does the territory actually look like? And --- most practically --- if \$35 billion and two decades of global effort have produced zero commercially viable quantum computers under the incumbent paradigm, are there alternative architectures that might actually deliver commercially useful computation?

This paper addresses four questions:

1. **What is "quantum," really?** Stripped of particle ontology, measurement

metaphysics, and anthropocentric baggage, what does the quantum formalism actually demand of a computational substrate?

1. **What is the solution space?** Which alternative computational paradigms ---

measurement-based, continuous-variable, topological, field-theoretic, analog --- are faithful to quantum reality, and how mature are they?

1. **What can actually be built?** Of these alternatives, which are commercially

manufacturable under real engineering, thermodynamic, and economic constraints?

1. **Is "quantum computing" even the right frame?** Or does the deeper lesson ---

that computation IS a physical process, and the substrate determines what can be computed efficiently --- point toward architectures where the physics does the computing directly, without being forced into a discrete symbolic scaffold?

## 2. What Is "Quantum," Really? An Ontological Stripping

Before we can ask what a "quantum computer" should look like, we must ask what "quantum" actually means --- not as a set of mathematical tools, but as a statement about what the world is made of and how it behaves.

### 2.1 The Textbook Picture and Its Projections

---

Most physicists learn quantum mechanics from textbooks that present it as a puzzle: "particles behave like waves, waves behave like particles, measurement collapses the wavefunction, and nobody knows why." This framing --- which Feynman called the "central mystery" --- is pedagogically useful but ontologically misleading. It takes classical categories (particle, wave, measurement, observer) as primitive and then expresses perplexity when quantum systems fail to conform to them.

The projection runs deep. We talk about "an electron" as if it were a tiny billiard ball with a position and momentum that we happen to be unable to measure simultaneously. We talk about "measurement" as if it were a passive observation that reveals pre-existing properties. We talk about "the observer" as if consciousness played a special role in physics. None of these are demanded by the formalism. All of them are classical intuitions we project onto a reality that operates differently.

### 2.2 The Minimal Ontological Commitments of Quantum Mechanics

---

If we strip away all classical projections --- particles, waves, measurement as revelation, observers as special --- what remains? What does the quantum formalism actually commit us to, ontologically?

The answer, distilled from seventy years of foundational work (relational quantum mechanics [Rovelli1996], QBism [Fuchs2014], consistent histories [Griffiths1984], and the decoherence program [Zurek2003]), is remarkably sparse:

**1. States are vectors in a complex Hilbert space.** A quantum system is described by a ray in a complex projective Hilbert space. This is the minimal mathematical structure needed to encode superposition, interference, and the Born rule. Nothing in this structure requires "particles," "waves," or "observers." It requires only a linear space with an inner product.

**2. Composite systems live in the tensor product, not the Cartesian product.** The state space of a composite system is NOT the set of all pairs of states of its components. It is a strictly larger space --- the tensor product --- that contains entangled states that cannot be factored. This is the formal statement of non-separability: the whole is not the sum of its parts.

**3. Properties are contextual.** The Kochen-Specker theorem [Kochen1967] proves that quantum systems cannot have pre-existing values for all observables simultaneously. A "property" of a quantum system is not an attribute of the system alone but of the system-in-measurement-context. Properties are relational, not intrinsic. This is not a limitation of our knowledge --- it is a structural fact about quantum reality.

**4. Evolution is unitary (when the system is isolated).** Isolated quantum systems evolve deterministically according to the Schrödinger equation. The "measurement problem" --- the apparent conflict between unitary evolution and wavefunction collapse --- is not a problem of physics but of insisting that measurement is something other than physical interaction between systems.

**5. Probabilities are Born-rule. Not ignorance probabilities.** The probabilities that appear in quantum mechanics are not epistemic (reflecting our ignorance of pre-existing values) but structural (reflecting the geometry of the Hilbert space). An electron in a superposition of spin-up and spin-down is not "really" spin-up or spin-down with us not knowing which. It is in a state that is genuinely neither, and the Born rule gives the probability of finding it in one or the other upon interaction with a measurement apparatus.

## 2.3 What Is NOT There: The Absent Ontology

---

Equally revealing is what these minimal commitments do NOT include:

- **No particles.** Nothing in the formalism requires that quantum systems be

localized in space, have trajectories, or persist as identifiable individuals. Quantum field theory, the more fundamental theory, explicitly rejects particle ontology: particles are excitations of fields, not fundamental entities.

- **No classical measurement apparatus.** The "Heisenberg cut" --- the boundary

between quantum system and classical measuring device --- is a pragmatic convenience, not an ontological boundary. Every measuring device is itself a quantum system. The appearance of classicality emerges from decoherence: the irreversible leakage of phase information into environmental degrees of freedom [Zurek2003].

- **No observers.** Consciousness plays no role in the formalism. The

"observer" is simply a physical system that becomes correlated with the measured system through interaction. QBism [Fuchs2014] correctly identifies that probabilities are agent-relative --- they encode an agent's expectations about future experiences --- but this does not make consciousness ontologically special. It makes probability theory agent-relative, which it already was.

- **No "collapse."** The projection postulate --- that measurement "collapses"

the wavefunction --- is not a dynamical law but a bookkeeping rule for updating probabilities conditional on new information. In decoherence-based interpretations, collapse is replaced by the physical process of environmental monitoring: the system becomes entangled with its environment, and the reduced density matrix of the system alone becomes effectively diagonal in the pointer basis. No collapse required [Schlosshauer2007].

## 2.4 The Relational Core: Quantum Mechanics as a Theory of Correlations

---

The positive ontological content that survives stripping is this: **\*\*quantum mechanics is a theory of correlations --- specifically, of the structure and dynamics of entanglement in composite quantum systems.\*\***

Consider what the quantum formalism is actually good at predicting: correlation statistics. The Bell inequalities, the CHSH game, quantum teleportation, superdense coding, quantum key distribution --- all of these are fundamentally about correlations between measurement outcomes on spatially separated subsystems. The "quantumness" lies not in the individual systems but in the correlation structure that connects them.

This is the insight of relational quantum mechanics [Rovelli1996]: the properties of a quantum system are not intrinsic but are relative to another system with which it interacts. There are no "states of the electron." There are only "states of the electron relative to the measurement apparatus." This is not idealism or anti-realism --- it is a structural claim about what kinds of facts quantum mechanics can encode.

From this perspective, the qubit-as-particle picture is not just empirically inadequate but categorically wrong. It treats quantum systems as objects with intrinsic properties that gates manipulate, when the actual quantum ontology is one of relational properties that emerge only through interaction. **\*\*The computational primitive is not the qubit**

but the correlation.\*\*

## 2.5 The Physical Meaning of "Quantum"

---

So what does "quantum" actually mean, physically? Three irreducible features:

**Discreteness of interaction.** The word "quantum" comes from the Latin *quantus* ("how much") and refers to the fact that certain physical quantities --- energy levels in bound systems, angular momentum, charge --- come in discrete units. This is not a statement about the nature of reality being "granular." It is a statement about the spectrum of certain operators being discrete, which follows from boundary conditions on the Schrödinger equation. Discreteness is an emergent property of confined systems, not a fundamental axiom.

**Superposition and interference.** The linearity of the Schrödinger equation means that if state A is possible and state B is possible, then any complex linear combination is also possible. This is the source of interference effects -- - the double-slit experiment, the Aharonov-Bohm effect --- and it has no classical analog. Superposition is not a statement about a particle "being in two places at once." It is a statement about the vector space structure of quantum states.

**Entanglement.** The tensor product structure of composite systems means that the state of a composite system is not determined by the states of its components. This is the most radical departure from classical physics and the one most directly relevant to computation. Entanglement is what makes quantum systems computationally interesting: the dimension of the state space grows exponentially with the number of components, but this exponential space is not freely accessible --- it is structured by the geometry of entanglement.

These three features --- discreteness, superposition, entanglement --- are the invariants. Everything else --- qubits, gates, circuits, measurements, error correction --- is scaffold.

## 3. The Solution Space: Computational Paradigms Faithful to Quantum Reality

If the qubit-gate-circuit scaffold is a map-territory confusion, what are the alternatives? What would a computational paradigm faithful to the ontological commitments of quantum mechanics actually look like? We survey five paradigms, ordered from most to least mature.

### 3.1 Measurement-Based Quantum Computing (MBQC)

---

**Core idea:** The computation IS the entanglement structure. Prepare a highly entangled resource state (a cluster state or graph state), then perform a sequence of single-qubit measurements whose outcomes --- and the choice of measurement bases, adapted to prior outcomes --- implement the computation. There are no gates. The "algorithm" is encoded in the pattern of measurements on the entangled substrate.

**Ontological fidelity:** High. In MBQC, entanglement is not a side effect or a resource to be managed --- it IS the computational primitive. The cluster state is a single, massively entangled state that encodes the entire computation. Gates are replaced by measurement patterns, which are inherently contextual (the basis choice at each step depends on prior outcomes). This aligns with the relational picture: computation proceeds through a sequence of contextual interactions (measurements) rather than through the manipulation of intrinsically-valued objects (qubits).

**Maturity:** MBQC is theoretically well-understood. Raussendorf and Briegel proved in 2001 that cluster-state MBQC is computationally universal [Raussendorf2001]. The framework has been extended to continuous-variable systems [Menicucci2006] and to fault-tolerant architectures. Experimental demonstrations exist at small scale (a few qubits in photonic systems).

**Limitations:** MBQC requires preparing large-scale entangled resource states, which is itself a major experimental challenge. The resource state must be prepared with high fidelity before computation begins, and any errors in the state propagate through the measurement sequence. However, topological cluster states --- which encode the resource state in a topologically protected subspace --- offer a path to fault tolerance that may be more natural than surface-code error correction in the circuit model.

**Commercial traction:** PsiQuantum's "fusion-based" quantum computing architecture is essentially an MBQC implementation using photonic qubits, where small entangled states are probabilistically fused into larger cluster states. This represents the most commercially advanced MBQC effort.

### 3.2 Continuous-Variable Quantum Computing (CV-QC)

---

**Core idea:** Replace discrete qubits --- two-level quantum systems --- with continuous quantum variables: the quadrature amplitudes of the electromagnetic field (position and momentum analogs). Qumodes --- the CV analog of qubits --- live in an infinite-dimensional Hilbert space, and computation proceeds through Gaussian operations (squeezing, displacement, beam-splitting) supplemented by non-Gaussian operations for universality.

**Ontological fidelity:** Very high. CV-QC treats the computational substrate as genuinely continuous --- a quantum field --- rather than discretizing it into artificial two-level systems. This is more faithful to the field-theoretic nature of quantum electrodynamics, where the fundamental entities are continuous field modes, not particles. The infinite-dimensional Hilbert space of a qumode encodes information in the continuous amplitude and phase of a field quadrature, which is closer to how quantum information is actually carried in physical systems (photons, phonons, collective excitations).

**Maturity:** CV-QC is theoretically well-developed but experimentally less mature than discrete-variable (qubit) approaches. Gottesman-Kitaev-Preskill (GKP) states provide a path to fault-tolerant CV computation by encoding discrete logical qubits in the continuous Hilbert space of a harmonic oscillator [Gottesman2001]. CV cluster states extend MBQC to the CV domain. Experimental demonstrations of CV entanglement, squeezing, and small-scale computation exist in optical systems.

**Commercial traction:** Xanadu is the leading commercial CV-QC company, developing photonic quantum processors based on squeezed-light qumodes. Their "Borealis" processor demonstrated Gaussian boson sampling at a scale that (arguably) exceeds classical simulability. The photonic platform is attractive because it operates at room temperature, integrates with existing fiber-optic infrastructure, and avoids the cryogenic requirements of superconducting qubits.

**The CV advantage:** Because CV systems encode information in continuous degrees of freedom, they can represent certain classes of functions --- Gaussian states and operations --- with exponentially fewer resources than discrete systems. This suggests that for problems with continuous structure (optimization over continuous variables, simulation of bosonic systems, machine learning with continuous data), CV architectures may have a natural advantage over qubit-based ones.

### 3.3 Topological Quantum Computing (TQC)

---

**Core idea:** Encode quantum information in globally-defined, topologically protected degrees of freedom --- specifically, the fusion and braiding properties of anyons (quasiparticle excitations in two-dimensional systems that are neither bosons nor fermions). Computation is performed by braiding anyons around each other; the result depends only on the topological class of the braid, not on the detailed trajectory, providing intrinsic protection against local noise.

**Ontological fidelity:** High, in a specific sense. TQC recognizes that quantum information is fundamentally about global, relational structure --- the topology of particle trajectories --- rather than about local states of individual particles. The "qubit" in TQC is not a physical two-level system but a topological degree of freedom that is distributed across multiple anyons. The information is stored non-locally, which is why it is protected from local perturbations.

This aligns with the relational picture: the computational DOF are relationships (braiding patterns), not objects (individual anyons). The topological protection arises precisely because the information is encoded in the global structure, not in any local degree of freedom.

**Maturity:** TQC is theoretically elegant but experimentally extremely challenging. The existence of non-abelian anyons --- the type needed for universal quantum computation --- was only experimentally confirmed in 2020 in fractional quantum Hall systems. Microsoft has invested heavily in Majorana zero modes (a type of non-abelian anyon) in semiconductor-superconductor heterostructures, but the experimental evidence has been contested and retracted [Frolov2021]. As of 2026, topological qubits have not been demonstrated at a level that would enable computation.

**The universality challenge:** The simplest non-abelian anyons (Ising anyons) support only a restricted set of gates (Clifford group) and require non-topological operations (magic state distillation) for universality. More exotic anyons (Fibonacci anyons) are universal but have not been experimentally realized. This means TQC may face the same error-correction overhead problem as the circuit model, just at a different level.

### 3.4 Quantum Field Computation: The Speculative Frontier

---

**Core idea:** If quantum field theory is the fundamental theory --- particles are excitations of fields, not fundamental entities --- then the ultimate computational substrate is not a collection of qubits but a quantum field. Computation would proceed by engineering the Hamiltonian of a quantum field such that its dynamics (ground state, time evolution, correlation functions) encode the solution to a computational problem.

**What this might look like:** A "quantum field computer" would be an analog device: a physical system --- perhaps an array of coupled superconducting resonators, a cold atom gas in an optical lattice, or a engineered metamaterial --- whose quantum field theory is designed so that measuring its correlation functions yields the answer to a hard computational problem. The "programming" would consist of shaping the Hamiltonian --- the interactions, couplings, and boundary conditions --- rather than specifying a sequence of gates.

**Ontological fidelity:** Maximal. This is the paradigm most faithful to quantum field theory, since it treats the field --- rather than its particle excitations --- as fundamental. The computational resource is the continuum of degrees of freedom of the field, and the computation is encoded in the field's correlation structure.

**Maturity:** Essentially zero. This is a speculative research direction. Some elements exist: quantum simulation (Feynman's original vision) can be seen as a primitive form of field computation, where one quantum system (the simulator) is designed to have the same Hamiltonian as another (the target system). But the idea of programming a quantum field theory as a general-purpose computer is unexplored.

**Why it might matter:** If the qubit is a scaffold and the field is the reality, then the most natural computational primitive is not the two-level system but the continuum field mode. The infinite-dimensional Hilbert space of a quantum field encodes vastly more information than a collection of qubits, and --- crucially --- it encodes it in a way that is physically natural, avoiding the artificial discretization that creates the error-correction overhead of the circuit model.

### 3.5 Analog Quantum Computing and Quantum Simulation

---

**Core idea:** Instead of building a universal digital quantum computer, build a special-purpose analog device whose natural dynamics solve a specific class of problems. Quantum simulation --- using one controllable quantum system to simulate another --- is the paradigm case. Quantum annealing (D-Wave's approach) is another: encode the problem in the energy landscape of a quantum system and let the system's natural relaxation dynamics find the ground state.

**Ontological fidelity:** Mixed. Analog approaches are faithful to the continuous, dynamical nature of quantum systems --- they let the physics do the computing rather than imposing a discrete symbolic scaffold. But they typically encode problems in a classical cost function (e.g., Ising Hamiltonian for optimization), which is itself a scaffold.

**Maturity:** D-Wave has deployed quantum annealers with thousands of qubits, and quantum simulators (cold atoms in optical lattices, trapped ions, Rydberg atom arrays) have simulated quantum many-body systems beyond the reach of classical computation. However, no analog quantum device has yet demonstrated a commercially relevant speedup over classical alternatives for a problem of practical interest.

**The fundamental limitation:** Analog quantum computers are not universal. Each device is designed for a specific class of problems, and the mapping between problem and device is itself a scaffold. For optimization problems, the question is whether quantum annealing offers an asymptotic advantage over classical heuristic solvers --- and the evidence to date suggests it does not [Ronnnow2014, Denchev2016].

## 4. What Can Actually Be Built? Commercially Manufacturable Paths

Theoretical fidelity to quantum ontology is necessary but not sufficient. A computational paradigm must also be physically realizable, scalable, and --- for commercial viability --- manufacturable under industrial constraints. We assess the leading physical platforms against these criteria.

### 4.1 Photonic Quantum Computing

---

**Platform:** Photons as qubits/qumodes; linear optical elements (beam splitters, phase shifters) for gates; single-photon detectors for measurement.

**Key players:** PsiQuantum (fusion-based MBQC with silicon photonics), Xanadu (CV squeezed-light qumodes), QuiX (integrated photonics).

**Advantages:**

- Room-temperature operation (no cryogenics).
- Existing semiconductor photonics manufacturing infrastructure.
- Natural fit for CV and MBQC paradigms.
- Photons are naturally decoherence-free for short propagation distances.
- Inherently networked: photons are the natural carrier of quantum information.

**Challenges:**

- Photon loss is the dominant error mechanism and scales exponentially with

circuit depth. While error correction can handle loss below a threshold, that threshold is demanding (~few percent per component).

- Single-photon sources and detectors are inefficient and noisy.
- Deterministic two-qubit gates are difficult in linear optics; most

architectures rely on probabilistic entanglement generation followed by multiplexing, which imposes enormous resource overhead.

- PsiQuantum's fusion-based approach requires millions of physical components

for a useful logical qubit, though they claim this is feasible in silicon photonics at wafer scale.

**Manufacturability verdict:** High. Silicon photonics leverages the existing semiconductor fabrication ecosystem. If the component-level performance thresholds can be met, photonic quantum computers could be manufactured at scale using existing fabs. This is the strongest commercial case in the quantum computing landscape as of 2026.

## 4.2 Neutral Atom Arrays

---

**Platform:** Individual neutral atoms (rubidium, strontium, ytterbium) trapped in optical tweezer arrays; Rydberg blockade for entangling gates; optical readout.

**Key players:** QuEra, Atom Computing, Pasqal.

**Advantages:**

- Atoms are identical (no fabrication variability).
- Long coherence times (seconds for hyperfine qubits).
- Reconfigurable: atoms can be rearranged during computation.
- Natural fit for analog quantum simulation (Rydberg atom arrays can simulate

Ising and XY models).

- Scaling to thousands of atoms demonstrated.

**Challenges:**

- Gate fidelities, while improving, are below fault-tolerance thresholds for many architectures (~99.5% for two-qubit gates, need >99.9%).
- Atomic rearrangement is slow (~milliseconds), limiting clock speed.
- Optical access for individual addressing becomes challenging at scale.
- Vacuum and laser infrastructure is bulky and expensive.

**Manufacturability verdict:** Medium. The components (lasers, vacuum chambers, optics) are commercial-off-the-shelf, and atom arrays scale well in principle. But the system integration complexity is high, and the cost per qubit is currently orders of magnitude above what would be needed for commercial deployment. This platform is most promising for specialized scientific applications (quantum simulation) in the near term.

## 4.3 Superconducting Qubits: The Incumbent

---

**Platform:** Josephson junction-based superconducting circuits (transmon, fluxonium, etc.) operated at millikelvin temperatures.

**Key players:** IBM, Google, Rigetti.

**Advantages:**

- Most mature platform by investment and engineering effort.
- Leverages semiconductor fabrication techniques.
- Fast gate speeds ( $\sim 10\text{-}100\text{ ns}$ ).
- Demonstrated error correction below threshold (Google's Willow, 2024).

#### Challenges:

- Cryogenic operation at  $\sim 10\text{ mK}$  requires dilution refrigerators, which are expensive, bulky, and power-hungry.

- Qubit coherence times are fundamentally limited by materials defects (two-level systems in amorphous dielectrics).

- Crosstalk between qubits increases with density.
- The error correction overhead ---  $\sim 1000$  physical qubits per logical qubit

for surface codes --- means that a useful fault-tolerant machine would require millions of physical qubits, each individually controlled.

- The thermodynamic cost of operating a million-qubit cryogenic system may be prohibitive.

**Manufacturability verdict:** Low for commercially useful scale. While small processors (100-1000 qubits) can be manufactured, the scaling requirements for fault-tolerant computation --- millions of physical qubits at millikelvin, each with individual control lines --- appear incompatible with commercial deployment at any plausible cost point. The superconducting platform may prove that fault-tolerant quantum computation is \*possible in principle\* while being commercially and thermodynamically impractical.

## 4.4 Trapped Ions

---

**Platform:** Individual atomic ions confined in electromagnetic traps; qubits encoded in electronic or hyperfine states; gates via laser or microwave pulses; readout via fluorescence.

**Key players:** IonQ, Quantinuum (Honeywell).

#### Advantages:

- Highest gate fidelities of any platform ( $>99.9\%$  for two-qubit gates).
- Long coherence times (seconds to minutes).
- All-to-all connectivity (ions can be shuttled).
- Quantinuum has demonstrated repeated error correction cycles.

#### Challenges:

- Gate speeds are slow ( $\sim \text{microseconds to milliseconds}$ ) compared to superconducting qubits.

- Scaling to large numbers of ions in a single trap is limited by motional mode crowding.

- Photonic interconnects for modular architectures add complexity and loss.
- System size, cost, and complexity scale poorly.

**Manufacturability verdict:** Low to medium. The exquisite control achieved in academic and industrial labs does not translate easily to high-volume manufacturing. Each ion trap is a precision instrument requiring ultra-high vacuum, stabilized lasers, and careful calibration. While trapped ions are excellent for demonstrating quantum algorithms at small scale, the path to commercial deployment at useful scale is unclear.

## 4.5 The Common Scaling Problem

---

Across all platforms, a common challenge emerges. To perform useful quantum computation --- solving a problem of genuine commercial value faster or cheaper than classical alternatives --- fault tolerance is required. Fault tolerance requires quantum error correction, which imposes an overhead of roughly  $10^2$  to  $10^3$  physical qubits per logical qubit. Useful computation requires perhaps 100-1000 logical qubits. This means a useful fault-tolerant quantum computer needs  $10^5$  to  $10^6$  physical qubits, each controlled with high fidelity.

No platform has demonstrated a credible path to manufacturing, cooling, controlling, and reading out a million high-fidelity qubits at a cost that would make commercial deployment viable. This is not merely an engineering challenge --- it may be a thermodynamic one. The entropy generated by controlling and reading out a million quantum degrees of freedom, multiplied by the error correction overhead, may exceed what can be dissipated in a practical device.

## 5. Post-Quantum: The Next Computational Frontier

If the qubit-gate-circuit paradigm is epistemically flawed and the alternative quantum paradigms face daunting engineering challenges, what comes next? Here we step back from "quantum computing" entirely and ask: what does the deeper lesson --- that computation IS a physical process and the substrate determines what can be computed --- imply for the future of computing?

### 5.1 The Fundamental Insight: Substrate IS Algorithm

---

The qubit delusion is a specific case of a broader pattern in computer science: the tendency to treat computation as abstract symbol manipulation, independent of the physical substrate that performs it. This abstraction --- the Church-Turing thesis, the von Neumann architecture, the distinction between hardware and software --- has been enormously productive. But it is, itself, a scaffold.

The deeper insight, which quantum computing inadvertently exposed, is that **the physical substrate IS part of the algorithm**. A quantum computer is not just a faster classical computer --- it accesses a different computational resource (entanglement, superposition) that has no classical analog. The "software-hardware" distinction breaks down when the physics of the hardware is the computational primitive.

This suggests that the next major advances in computing will come not from making faster transistors but from finding physical substrates whose natural dynamics solve computationally hard problems directly. The substrate IS the algorithm.

### 5.2 Thermodynamic Computing

---

**Core idea:** Treat computation as a thermodynamic process. Physical systems evolve toward thermal equilibrium by minimizing a free energy functional. If we can encode a computational problem as the free energy landscape of a physical system, the system's natural relaxation dynamics will "compute" the solution by finding the free energy minimum.

This is not a new idea --- it is how nature "computes" protein folding, crystal growth, phase separation, and neural dynamics. What is new is the possibility of engineering thermodynamic systems specifically for computation.

#### Physical implementations:

- **Ising machines:** Arrays of coupled oscillators (optical parametric oscillators, CMOS LC oscillators) that naturally minimize an Ising Hamiltonian. Coherent Ising machines [Marandi2014] have demonstrated solving combinatorial optimization problems orders of magnitude faster than classical heuristics for certain problem classes --- not by being "quantum" but by being analog and massively parallel.

- **Probabilistic bits (p-bits):** Stochastic nanomagnetic devices that fluctuate between states with probabilities governed by a Boltzmann distribution. Networks of p-bits can perform probabilistic inference, optimization, and sampling by exploiting thermal fluctuations rather than fighting them [Camsari2017].

- **Memristive neural networks:** Crossbar arrays of memristive devices that perform matrix-vector multiplication in the analog domain via Ohm's law and Kirchhoff's current law. This is not quantum --- it's classical physics --- but it leverages the physical substrate (resistive switching, charge transport) for computation directly.

**Commercial traction:** Several companies are developing Ising machines and p-bit architectures for optimization and machine learning. These are not "quantum computers" in the gate-model sense, but they exploit physical dynamics --- thermal fluctuations, phase transitions, collective modes --- for computation. They do not require cryogenics, error correction, or isolation from the environment. In fact, they USE the environment (thermal bath) as a computational resource.

### 5.3 Neuromorphic Computing

---

**Core idea:** Build computers that operate on the same principles as biological neural systems: massively parallel, asynchronous, analog, event-driven, and co-locating memory and computation.

#### Key principles:

- **Spike-based coding:** Information is encoded in the timing of discrete electrical pulses (spikes), not in continuous voltage levels. This is energy-efficient because spikes consume power only when they occur.

- **Event-driven computation:** Neurons compute only when they receive input, not on a global clock cycle. This eliminates the wasted energy of clock distribution in synchronous digital systems.

- **In-memory computation:** Synaptic weights are stored and updated at the point of computation (the synapse), eliminating the von Neumann bottleneck of shuttling data between memory and processor.

- **Stochastic dynamics:** Neural systems exploit noise --- synaptic stochasticity, ion channel fluctuations --- for exploration, learning, and robustness.

#### Physical implementations:

- **CMOS neuromorphic chips:** Intel's Loihi 2, IBM's TrueNorth, and

SpiNNaker use conventional CMOS technology with neuromorphic architectures. Loihi 2 achieves  $\sim 15$  pJ per synaptic operation, compared to  $\sim 100$  pJ for equivalent operations on conventional processors.

- **Memristive neuromorphic systems:** Memristors --- resistors with memory ---

can implement synapses directly in analog hardware, performing matrix-vector multiplication with  $O(1)$  energy per operation regardless of matrix size.

- **Photonic neuromorphic systems:** Optical neural networks perform linear

operations (matrix multiplication) through free-space or integrated photonic interference, at the speed of light and with minimal energy dissipation.

**Why this matters:** Neuromorphic computing is not "quantum computing lite." It represents a genuinely different computational paradigm --- one where the physical substrate (analog electronics, photonics, memristive materials) performs the computation directly, rather than being forced into a digital, clocked, von Neumann scaffold. The  $\sim 10^4\times$  energy efficiency advantage of biological neural systems over conventional computers is not a quirk of evolution --- it reflects the efficiency of analog, event-driven, in-memory computation.

## 5.4 Optical Computing (Classical)

---

**Core idea:** Use light --- photons propagating through free space, fibers, or integrated waveguides --- to perform computation. Linear optical elements (lenses, beam splitters, spatial light modulators) perform Fourier transforms, convolutions, and matrix multiplications at the speed of light, with zero energy dissipation in the optical path.

**Key principles:**

- **Fourier optics:** A lens performs a spatial Fourier transform --- an  $O(N \log N)$

operation in digital --- in a single pass of light, at the speed of light, with zero computational energy.

- **Interference-based computation:** Interference between coherent optical

fields naturally computes inner products, the fundamental operation of linear algebra.

- **Wavelength-division multiplexing:** Different wavelengths can carry

independent computational channels through the same physical medium.

**Commercial traction:** Several startups (Lightmatter, Lightelligence, Optalysys) are developing optical processors for matrix multiplication and convolution --- the dominant operations in deep learning. These are not "quantum" --- they use classical (coherent) light --- but they exploit the physical properties of light for computation.

Optical matrix multipliers can, in principle, perform matrix-vector products with  $O(1)$  energy scaling in the optical path (the energy cost is dominated by the input/output conversion, not the computation itself). For sufficiently large matrices ( $>1000 \times 1000$ ), optical processors may offer orders-of-magnitude energy advantages over digital electronics.

## 5.5 What Unifies These Paradigms

---

Thermodynamic, neuromorphic, and optical computing share a common philosophical core: **let the physics do the computing**. Instead of forcing computation into a universal digital scaffold (logic gates, clock cycles, von Neumann architecture), find a physical system whose natural dynamics solve the problem directly.

This is the same insight that motivated Feynman's original proposal for quantum simulation --- "let the quantum system simulate itself" --- but applied more broadly. A quantum computer is a special case: a physical system whose (substrate, Hamiltonian, dynamics) are engineered to perform a specific class of computations that exploit quantum superposition and entanglement. But there are many other physical systems --- thermal baths, optical fields, neural networks, chemical reaction networks --- whose natural dynamics solve computationally hard problems.

The future of computing, we suggest, is not a "quantum computer" in the qubit-gate sense but a **physical computation ecosystem**: a diverse set of special-purpose devices --- thermodynamic optimizers, neuromorphic inference engines, optical matrix processors, analog quantum simulators --- each exploiting the physics of its substrate for the class of problems it is naturally suited to solve.

## 6. Principles for Post-Particle Computation

From the analysis above, we distill five principles that should guide the development of post-qubit computational architectures:

**Principle 1: The substrate IS the algorithm.** Choose physical substrates whose natural dynamics compute the target problem class directly, rather than forcing generic substrates into universal computation.

**Principle 2: Correlation over particle.** Encode information in correlation structures (entanglement, phase coherence, collective modes) rather than in local states of individual degrees of freedom.

**Principle 3: Noise as resource, not enemy.** Exploit thermal fluctuations, stochastic dynamics, and environmental coupling for exploration, error tolerance, and energy efficiency, rather than fighting them with active error correction.

**Principle 4: Analog where natural, digital where necessary.** Use analog physical dynamics for the computationally hard part (optimization, sampling, matrix multiplication) and digital control for programmability and reconfigurability.

**Principle 5: Manufacturability as first-class constraint.** A computational paradigm that requires a million cryogenic control lines is not commercially viable regardless of its theoretical elegance. Design for the manufacturing ecosystem that actually exists.

## 7. The Research Agenda, Operationalized

We close with specific, falsifiable research directions that operationalize these principles. Each is tagged with feasibility [near-term: 1-3 years], [medium-term: 3-7 years], or [speculative: 7+ years].

### 7.1 Thermodynamic Optimization Engines [near-term]

---

Build and benchmark Ising machines, p-bit networks, and coupled oscillator arrays on practically relevant optimization problems (supply chain optimization, portfolio allocation, molecular docking). The key metric is not "quantum advantage" but "thermodynamic advantage": does the physical system find better solutions faster than classical heuristics at lower energy cost?

**Falsification criterion:** If thermodynamic optimizers consistently fail to outperform classical heuristics (simulated annealing, genetic algorithms) on real-world optimization instances, the paradigm is invalid.

## 7.2 Entanglement-First Quantum Architectures [medium-term]

---

Design quantum processors where the computational primitive is the preparation and manipulation of entanglement structure, not individual qubits. This means:

- MBQC architectures with cluster states as the fundamental resource.
- CV architectures where Gaussian entanglement is the computational substrate.
- Architectures where the physical DOF is a continuum field mode, not a two-level system.

**Falsification criterion:** If an MBQC or CV architecture consistently underperforms a circuit-model architecture with the same physical resources, the paradigm is not superior.

## 7.3 Neuromorphic Inference at Scale [near-term]

---

Deploy neuromorphic systems (CMOS or memristive) for large-scale inference in deep learning, where the energy advantage of analog, in-memory computation is most pronounced. The target: a  $10^3\times$  energy reduction for transformer inference compared to GPU-based systems, validated on standard benchmarks.

**Falsification criterion:** If neuromorphic systems cannot achieve  $>10\times$  energy advantage on at least one commercially relevant inference workload, the paradigm is not commercially viable.

## 7.4 Optical Linear Algebra Accelerators [near-term]

---

Develop integrated photonic processors for matrix multiplication and convolution at scale ( $>1000\times 1000$  matrices). The target: demonstrating a  $10^2\times$  energy-per-operation advantage over digital electronics at a scale relevant to deep learning inference.

**Falsification criterion:** If optical matrix multipliers cannot maintain an accuracy-appropriate precision ( $>4$  bits effective) at scale while achieving  $>10\times$  energy advantage, the paradigm is not viable.

## 7.5 Field-Theoretic Computation [speculative]

---

Investigate whether quantum field theories --- lattice gauge theories, conformal field theories, tensor network representations --- can serve as computational substrates. The question: can we design a physical system whose Hamiltonian is the QFT of a computationally hard problem, and whose correlation functions directly encode the solution?

**Falsification criterion:** If no mapping between a computationally hard problem class and a physically realizable QFT is found within five years of serious investigation, the paradigm should be downgraded to "speculative foundations research."

## 8. Conclusion

This paper has argued that the diagnosis offered by "The Qubit Delusion" implies a constructive research program that extends beyond merely fixing the qubit-gate-circuit model. The deeper lesson is that computation is an inherently physical process: the substrate matters, and the choice of computational scaffold is never ontologically neutral.

The qubit-gate-circuit scaffold projects particle ontology onto a relational reality. The error-correction paradigm fights decoherence rather than harnessing it. The commercial quantum computing industry has optimized for narrative production rather than physical fidelity. These are symptoms of a deeper pattern: the tendency to treat computation as abstract symbol manipulation, independent of the physical world that performs it.

The constructive alternative is not a single new computational paradigm but a **physical computation ecosystem**: a diverse portfolio of special-purpose devices --- thermodynamic optimizers, neuromorphic inference engines, optical matrix processors, entanglement-first quantum simulators, and ultimately field-theoretic computers --- each matched to the problem class that its physical substrate naturally solves.

This ecosystem will not emerge from the incumbent quantum computing industry, whose institutional incentives reward the appearance of progress over falsifiable results. It will emerge from research programs that take seriously the ontological commitments of quantum mechanics and the physical constraints of actual manufacturing --- programs that ask not "how do we build a better qubit?" but "what physical system, if any, can solve this problem faster, cheaper, or more energy-efficiently than any classical alternative?"

The qubit delusion is not the end of quantum computing. It is the beginning of honest computation --- computation that respects what physics actually says about the world.

---

## Acknowledgments

This work builds on the foundational insights of Rovelli (relational QM), Fuchs (QBism), Zurek (quantum Darwinism), Raussendorf & Briegel (MBQC), and Feynman (quantum simulation as physical computation).

## References

- @Rovelli1996: C. Rovelli, "Relational Quantum Mechanics," Int. J. Theor. Phys. 35, 1637 (1996).
- @Fuchs2014: C. A. Fuchs et al., "An Introduction to QBism with an Application to the Locality of Quantum Mechanics," Am. J. Phys. 82, 749 (2014).
- @Griffiths1984: R. B. Griffiths, "Consistent Histories and the Interpretation of Quantum Mechanics," J. Stat. Phys. 36, 219 (1984).
- @Zurek2003: W. H. Zurek, "Decoherence, Einselection, and the Quantum Origins of the Classical," Rev. Mod. Phys. 75, 715 (2003).
- @Kochen1967: S. Kochen and E. P. Specker, "The Problem of Hidden Variables in Quantum Mechanics," J. Math. Mech. 17, 59 (1967).
- @Schlosshauer2007: M. Schlosshauer, "Decoherence and the Quantum-to-Classical Transition" (Springer, 2007).
- @Raussendorf2001: R. Raussendorf and H. J. Briegel, "A One-Way Quantum Computer," Phys. Rev. Lett. 86, 5188 (2001).
- @Menicucci2006: N. C. Menicucci et al., "Universal Quantum Computation with Continuous-Variable Cluster States," Phys. Rev. Lett. 97, 110501 (2006).
- @Gottesman2001: D. Gottesman et al., "Encoding a Qubit in an Oscillator," Phys. Rev. A 64, 012310 (2001).
- @Frolov2021: S. Frolov, "Majorana Fiasco," Nature 591, 526 (2021).
- @Ronnow2014: T. F. Ronnow et al., "Defining and Detecting Quantum Speedup," Science 345, 420 (2014).
- @Danchev2016: V. S. Danchev et al., "What is the Computational Value of Finite-Range Tunneling?" Phys. Rev. X 6, 031015 (2016).

- @Marandi2014: A. Marandi et al., "Network of Time-Multiplexed Optical Parametric Oscillators as a Coherent Ising Machine," Nat. Photonics 8, 937 (2014).
- @Camsari2017: K. Y. Camsari et al., "Stochastic p-Bits for Invertible Logic," Phys. Rev. X 7, 031014 (2017).

## Changelog

---

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