

The Problem-Substrate Mapping: A Framework for Honest Computational Investment

Rowan Brad Quni-Gudzin

2026-07-08

Abstract. Three papers have established that the qubit-gate-circuit model is an epistemic failure, that alternative paradigms exist with greater ontological fidelity, and that fundamental physical limits impose honest boundaries on what any computational paradigm can deliver. This paper turns from critique to portfolio: given everything we now know, what should we actually build? We propose a systematic framework for matching computational problem classes to optimal physical substrates --- the Problem-Substrate Mapping --- and derive concrete investment theses for near-term (1-3 year), medium-term (3-7 year), and long-term (7-15 year) horizons. For each problem class (optimization, linear algebra, probabilistic inference, quantum simulation, cryptography, general-purpose computation), we identify the physical substrate that minimizes joules-per-solution at commercially relevant scale and assess its current technology readiness level. The resulting portfolio allocates roughly 40% to thermodynamic/analog computing, 25% to photonic/optical, 15% to neuromorphic, 10% to analog quantum simulation, 5% to reversible classical, and 5% to fault-tolerant quantum --- a dramatic departure from the ~90% allocation to gate-model quantum computing that characterizes current public and private investment.

Keywords. - computational investment, - problem-substrate mapping, - thermodynamic computing, - neuromorphic computing, - optical computing, - quantum computing, - technology portfolio, - R&D strategy

1. Introduction: From Critique to Portfolio

The first three papers in this series have established a foundation:

- **Phase I** ("The Qubit Delusion"): The qubit-gate-circuit model is an

epistemic failure --- a projection of particle ontology onto a relational, field-theoretic reality. The \$35 billion quantum computing industry has produced zero commercially viable machines because it has been optimizing the wrong scaffold.

- **Phase II** ("Beyond the Qubit"): Alternative paradigms --- measurement-based,

continuous-variable, topological, field-theoretic, thermodynamic, neuromorphic, optical --- exist with greater ontological fidelity and, in many cases, better commercial manufacturability.

- **Phase III** ("The Physics of Computation"): Fundamental physical limits ---

Landauer, Margolus-Levitin, Bremermann, Bekenstein --- define honest boundaries within which any computational paradigm must operate. Quantum error correction multiplies the thermodynamic cost of computation by 10^2 to 10^3 , meaning that only exponential algorithmic speedups can overcome the joules-per-solution penalty.

This paper turns from critique to construction. Given everything we now know --- about the epistemic failure of the qubit model, about the landscape of alternatives, about the fundamental physical limits --- what should we actually *build*? What should investors fund? What should government research agencies prioritize? What should entrepreneurs bet their careers on?

The answer is not a single technology. It is a **portfolio** --- a diversified allocation of intellectual and financial capital across multiple computational substrates, each matched to the problem class that its natural physics most efficiently solves.

We call this framework the **Problem-Substrate Mapping (PSM)**. It consists of:

1. A taxonomy of commercially relevant computational problem classes.

2. For each class, a mapping to the physical substrate(s) that minimize

joules-per-solution at commercially relevant scale.

1. A technology readiness assessment for each substrate-problem pair.

2. A recommended investment allocation across near-term, medium-term, and

long-term horizons.

1. Falsifiable milestones for each allocation.

2. Problem Classes and Their Physical Signatures

Every computational problem has a physical signature --- a pattern of information flow, memory access, arithmetic intensity, and parallelism that determines which physical substrate can solve it most efficiently. We identify six commercially relevant problem classes.

2.1 Optimization

What it is: Finding the minimum (or maximum) of a cost function over a discrete or continuous domain. Examples: supply chain optimization, portfolio allocation, vehicle routing, chip placement, protein folding, training neural networks (gradient descent is optimization).

Physical signature: The problem is naturally expressed as energy minimization. The cost function IS a Hamiltonian; the solution IS the ground state. Optimization problems are fundamentally thermodynamic: they ask "what is the lowest-energy configuration of this system?"

Computational demands: Exploration of a rugged energy landscape. The challenge is escaping local minima to find the global minimum. Classical heuristics (simulated annealing, genetic algorithms, gradient descent with momentum) already exploit thermal fluctuations as an exploration mechanism.

Natural substrate match: Physical systems that natively minimize free energy --- Ising machines, coupled oscillators, memristive crossbar arrays, and (potentially) quantum annealers. The physics does the optimization directly: the system evolves toward its ground state, and reading out that state gives the solution.

2.2 Linear Algebra

What it is: Matrix multiplication, singular value decomposition, eigenvalue computation, linear system solving. These operations dominate scientific computing, machine learning (every transformer forward pass is a sequence of matrix multiplications), and signal processing.

Physical signature: Linear algebra is fundamentally about inner products --- the multiplication and summation of vectors. This is the same operation that physical interference performs: when two coherent waves overlap, their amplitudes add, and the intensity encodes the inner product.

Computational demands: High arithmetic intensity ($O(N^3)$ for matrix multiply). Memory bandwidth is the bottleneck on conventional architectures. The computation is highly regular and parallelizable.

Natural substrate match: Optical processors. A lens performs a Fourier transform --- an $O(N \log N)$ linear operation --- in a single pass of light at zero computational energy. Integrated photonic circuits can perform matrix multiplication through cascaded Mach-Zehnder interferometers [Shen2017]. The energy cost is dominated by

input/output conversion (electrical to optical and back), not by the computation itself.

2.3 Probabilistic Inference

What it is: Computing conditional probabilities, sampling from complex distributions, Bayesian updating, graphical model inference, generative modeling. These operations are central to machine learning, risk assessment, decision theory, and scientific data analysis.

Physical signature: Probabilistic inference is naturally expressed as sampling from a Boltzmann distribution --- the same distribution that physical systems at thermal equilibrium naturally occupy. The problem asks "what is the most probable configuration given the evidence?" --- which is isomorphic to "what is the lowest-energy configuration given the constraints?"

Computational demands: Sampling from high-dimensional distributions is the computational bottleneck. Markov Chain Monte Carlo (MCMC) is the workhorse, but it mixes slowly for complex distributions. The challenge is efficient exploration of probability space.

Natural substrate match: Probabilistic bits (p-bits) --- nanomagnetic or CMOS devices that fluctuate between 0 and 1 with probabilities governed by a tunable energy landscape [Camsari2017]. Networks of p-bits naturally perform Boltzmann sampling. Neuromorphic processors also excel at probabilistic inference through spike-based stochastic computation.

2.4 Quantum Simulation

What it is: Simulating the behavior of quantum many-body systems --- molecules, materials, nuclear matter, quantum fields --- that are exponentially hard to simulate on classical computers due to the exponential growth of the Hilbert space.

Physical signature: The problem IS a quantum system. The Hamiltonian of the target system is the same mathematical object as the Hamiltonian of a controllable quantum device. This is Feynman's original insight [Feynman1982]: let the quantum system simulate itself.

Computational demands: Exponential classical complexity. The wavefunction of N interacting quantum particles requires $O(\exp(N))$ classical bits to represent. No classical computer --- reversible or otherwise --- can overcome this exponential scaling.

Natural substrate match: Analog quantum simulators --- cold atoms in optical lattices, trapped ion arrays, Rydberg atom arrays, superconducting circuits --- where the physical Hamiltonian is engineered to match the target Hamiltonian. The system evolves under its natural dynamics, and measurement of correlation functions yields the quantities of interest.

This is the one problem class where quantum physics provides a genuine, in-principle exponential advantage --- and it does not require fault tolerance, error correction, or universal gate sets. It requires only that the simulator is sufficiently coherent and controllable to faithfully reproduce the target Hamiltonian's physics. This is a far lower bar than fault-tolerant universal quantum computation.

2.5 Cryptography and Number Theory

What it is: Factoring large integers, computing discrete logarithms, and related number-theoretic problems that underpin public-key cryptography (RSA, ECC). Shor's algorithm provides an exponential quantum speedup for these problems.

Physical signature: These problems have no natural physical analog. They require the kind of coherent quantum interference that only a universal fault-tolerant quantum computer can provide. The physical substrate must support the quantum Fourier transform --- the core subroutine of Shor's algorithm --- at a scale and fidelity far beyond current capability.

Computational demands: For 2048-bit RSA, Shor's algorithm requires approximately 4,100 logical qubits and 10^9 Toffoli gates [Gheorghiu2019]. With surface-code error correction at a physical error rate of 10^{-3} , this translates to $\sim 10^7$ physical qubits and $\sim 10^{11}$ physical operations. The joules-per-solution analysis from Phase III suggests this is thermodynamically possible but commercially distant --- the cryogenic and error-correction overhead is enormous.

Natural substrate match: Fault-tolerant universal quantum computers --- the very paradigm that Phases I-III critique. For this specific problem class, the critique does not apply: the exponential algorithmic speedup can, in principle, overcome the thermodynamic overhead. The question is whether we can build a device of sufficient scale within any commercially relevant timeframe. The answer, as of 2026, is: not in the next 15 years.

2.6 General-Purpose Sequential Computation

What it is: The kind of computation that dominates the global compute fleet: operating systems, databases, web servers, business logic, compilers, video games, user interfaces. Code with branches, loops, function calls, pointer chasing, and irregular memory access patterns.

Physical signature: Highly sequential, branch-heavy, memory-intensive. The von Neumann architecture is not an arbitrary convention --- it reflects the structure of the problems being solved. General-purpose computation resists parallelization and resists analog implementation because its control flow is fundamentally discrete and conditional.

Computational demands: Low arithmetic intensity, high memory bandwidth, unpredictable branches. The bottleneck is not floating-point throughput but the memory wall --- the growing gap between processor speed and memory access time.

Natural substrate match: Reversible classical CMOS operating near the Landauer limit for energy-efficient sequential computation; conventional CMOS for everything else. Neuromorphic and optical processors are poor fits for this problem class because they are optimized for regular, parallel, high-arithmetic-intensity workloads. Quantum computers are useless for it.

3. The Substrate Portfolio

Based on the problem-substrate mapping above, we can now construct a concrete portfolio of computational substrates, each allocated to the problem class it most naturally solves.

3.1 Ising Machines and Thermodynamic Solvers

Problem class: Optimization.

How it works: An array of coupled oscillators --- optical parametric oscillators, CMOS LC tanks, or nanomagnetic spin systems --- is configured so that the system's energy landscape encodes the optimization problem's cost function. The system is allowed to relax toward its ground state through natural dissipative dynamics. The final configuration is read out as the solution.

Technology readiness: Coherent Ising machines have demonstrated solving MAX-CUT problems with thousands of spins on optical platforms [Honjo2021]. CMOS-based Ising solvers (Hitachi, Fujitsu, Toshiba) are commercially available for combinatorial optimization at the 1,000-100,000 variable scale. These are not research prototypes ---

they are shipping products.

Joules-per-solution advantage: For sufficiently large optimization problems ($N > 1,000$), Ising machines can achieve 10^1 to $10^3\times$ energy advantage over classical heuristics running on conventional processors. The advantage comes from massive parallelism (all spins update simultaneously) and the elimination of the memory wall (computation and "memory" are the same physical system).

Near-term investment thesis (1-3 years): Deploy Ising machines for real-world optimization in logistics, finance, and manufacturing. The technology is mature enough for commercial deployment. The limiting factor is not hardware capability but problem mapping --- encoding real optimization problems into Ising form.

3.2 Optical Processors

Problem class: Linear algebra (matrix multiply, convolution).

How it works: Coherent light propagates through an array of programmable beam splitters and phase shifters implemented in silicon photonics. The interference pattern at the output encodes the matrix-vector product of the input vector with the matrix encoded in the photonic circuit.

Technology readiness: Integrated photonic matrix multipliers at the 64×64 scale have been demonstrated [Shen2017]. Scaling to $1,000\times 1,000$ is expected within 2-3 years. The manufacturing infrastructure exists: silicon photonics leverages the same fabs that produce CMOS electronics.

Joules-per-solution advantage: For matrix multiplication at scale ($N > 1,000$), optical processors can achieve 10^2 to $10^3\times$ energy advantage over GPUs. The optical path dissipates essentially zero energy; the energy cost is dominated by laser power and photodetection. Unlike electronic processors, the energy per operation does NOT scale with matrix size --- the light does the computation "for free."

Near-term investment thesis (1-3 years): Deploy optical processors as inference accelerators for large neural networks, where matrix multiplication dominates runtime and energy consumption. Companies: Lightmatter, Lightelligence, Optalysys.

3.3 Neuromorphic and p-Bit Processors

Problem class: Probabilistic inference, pattern recognition, low-power sensing.

How it works: Spiking neural networks implemented in CMOS (Intel Loihi, IBM TrueNorth) or memristive crossbar arrays perform computation through the timing of discrete electrical pulses rather than continuous voltage levels. p-bits --- stochastic nanomagnetic devices --- naturally sample from Boltzmann distributions for probabilistic inference.

Technology readiness: Loihi 2 is commercially available and has demonstrated $\sim 10^3\times$ energy advantage over GPUs for specific inference workloads [Davies2018]. Memristive neuromorphic systems remain at the research prototype stage but have demonstrated proof-of-concept matrix multiplication at ~ 10 fJ per operation.

Joules-per-solution advantage: For inference workloads (the dominant cost in deployed AI), neuromorphic processors achieve 10^2 to $10^3\times$ energy advantage over GPUs. For probabilistic sampling, p-bit networks can achieve similar advantages over classical MCMC.

Near-term investment thesis (1-3 years): Deploy neuromorphic processors for edge AI --- always-on sensing, keyword spotting, anomaly detection --- where the energy budget is severely constrained (microwatts to milliwatts). Data center deployment for large-scale inference will follow as the technology matures.

3.4 Analog Quantum Simulators

Problem class: Quantum simulation (many-body physics, quantum chemistry, materials science).

How it works: A controllable quantum system --- cold atoms in an optical lattice, trapped ions, Rydberg atom arrays, or superconducting circuits --- is engineered to have the same Hamiltonian as the target quantum system. The simulator evolves under its natural dynamics, and measurements of correlation functions yield the quantities of interest. No error correction is required because the computation IS the physical evolution --- the system does not need to maintain a logical qubit; it only needs to be sufficiently coherent to faithfully reproduce the target physics.

Technology readiness: Cold atom quantum simulators have simulated the Fermi-Hubbard model at scales (~100 sites) that challenge classical simulation [Mazurenko2017]. Rydberg atom arrays have probed quantum phase transitions and non-equilibrium dynamics in Ising-like systems with hundreds of atoms. These are research demonstrations, not commercial products, but the path to useful quantum simulation is far shorter than the path to fault-tolerant quantum computation.

Joules-per-solution advantage: For quantum simulation problems at sufficient scale ($N > 50$ strongly interacting particles), analog quantum simulators may already achieve joules-per-solution advantage over classical simulation. The crossover point depends on the specific problem and the classical competitor (exact diagonalization vs. tensor networks vs. quantum Monte Carlo).

Medium-term investment thesis (3-7 years): Fund analog quantum simulation as the primary quantum computing research program. The goal is not a universal quantum computer but a suite of special-purpose simulators for the most commercially valuable quantum simulation problems: catalyst design, battery materials, pharmaceutical molecular dynamics.

3.5 Reversible Classical Computing

Problem class: General-purpose computation at ultra-low energy.

How it works: Classical logic gates are operated adiabatically --- slowly enough that the energy used to charge a capacitor is recovered when it is discharged, rather than being dissipated as heat. Information is never erased except at final readout, so the Landauer bound is paid only once per computation, not once per operation.

Technology readiness: Adiabatic microprocessors have been demonstrated with energy dissipation approaching 1% of the Landauer limit --- approximately 0.03 kT per operation [Snider2012]. These are laboratory demonstrations with simple circuits, not commercial products, but the physics is sound.

Joules-per-solution advantage: For general-purpose computation, reversible processors could, in principle, achieve $10^4\times$ energy advantage over conventional CMOS. In practice, the overhead of reversible logic (reverse computation for uncomputation, additional control circuitry) may reduce this to 10^1 to $10^2\times$. Still --- a $10\times$ to $100\times$ improvement in the energy efficiency of general-purpose computation would be transformative.

Long-term investment thesis (7-15 years): Fund fundamental research in reversible and adiabatic computing as the long-term path to energy-efficient general-purpose computation. This is not a near-term commercial play --- the market does not demand it because conventional CMOS still has decades of efficiency scaling ahead. But as CMOS approaches fundamental limits, reversible computing will become essential.

3.6 Fault-Tolerant Quantum Computing

Problem class: Cryptography (factoring, discrete log), and possibly quantum simulation at scales beyond analog capability.

How it works: Universal gate-model quantum computing with quantum error correction --- the paradigm that Phases I-III critique. The critique stands: this is the most ontologically unfaithful, thermodynamically expensive, and commercially distant computational paradigm. But for one problem class --- cryptography --- it may be the only path.

Technology readiness: No fault-tolerant quantum computer exists. Google's Willow processor (2024) demonstrated error correction below the surface-code threshold --- a genuine scientific achievement --- but at a scale (105 qubits) that is $10^5\times$ smaller than what is needed for useful computation.

Joules-per-solution advantage: Potentially enormous for factoring --- if a fault-tolerant quantum computer can be built. The joules-per-solution crossover for Shor's algorithm on 2048-bit RSA is estimated at $\sim 10^7$ physical qubits, requiring a cryogenic infrastructure of unprecedented scale. The thermodynamic analysis from Phase III suggests this is possible in principle but commercially distant.

Long-term investment thesis (7-15+ years): Maintain a small, rigorously-evaluated research program in fault-tolerant quantum computing, funded primarily through government research agencies with strong independent verification requirements. Private venture capital should NOT fund fault-tolerant quantum computing: the timeline-to-revenue is incompatible with VC fund horizons, and the information asymmetry between company claims and investor understanding creates an adverse selection problem.

4. The Investment Portfolio

We can now propose a concrete allocation of research and investment capital across substrates and time horizons.

4.1 Near-Term (1-3 Years)

Substrate	Allocation	Rationale	Measurable Milestone
Ising/thermodynamic solvers	25%	Commercially deployable for optimization; low technical risk	Solve $N>10,000$ variable real-world logistics problem at lower cost than classical
Optical processors	20%	Silicon photonics manufacturability; large inference market	Demonstrate $100\times$ energy advantage for transformer inference at batch=1
Neuromorphic processors	20%	Proven efficiency for edge AI; commercial products exist	Deploy in >10 consumer devices at <1 mW always-on power
Analog quantum simulation	15%	Nearest path to genuine quantum advantage; high scientific value	Simulate a quantum system beyond exact classical diagonalization
p-bit probabilistic networks	10%	Emerging; high potential for inference and optimization	Demonstrate Boltzmann sampling at $>10\times$ energy advantage vs MCMC
Conventional CMOS optimization	10%	Still dominates; algorithmic innovations matter	---

4.2 Medium-Term (3-7 Years)

Substrate	Allocation	Rationale	Measurable Milestone
Analog quantum simulation	30%	Scale from 100 to 10,000 atoms; target materials/pharma	Simulate catalyst reaction pathway at chemical accuracy
Optical processors	25%	Scale from 64×64 to 10,000×10,000 photonic circuits	Replace GPU cluster for inference in production data center
Neuromorphic processors	20%	Scale from edge to data center; memristive integration	Memristive crossbar at 1,000×1,000 scale in commercial product
Reversible/adiabatic CMOS	15%	Foundational research; prepare for post-CMOS era	Demonstrate reversible processor at 1 MHz, 1% Landauer limit
p-bit networks	10%	Scale to 10 ⁶ p-bits; target combinatorial optimization	Solve TSP at N>1,000 with joules-per-solution advantage

4.3 Long-Term (7-15+ Years)

Substrate	Allocation	Rationale	Measurable Milestone
Reversible/adiabatic CMOS	30%	Path to Landauer-limit general-purpose computing	General-purpose reversible processor at commercial scale
Photonic quantum (MBQC, CV)	25%	Room-temperature quantum; avoids cryogenic overhead	Logical qubit with error rate <10 ⁻⁶ at room temperature
Analog quantum simulation	20%	Full-scale materials and drug design	New catalyst or drug candidate discovered via quantum simulation
Fault-tolerant QC (crypto)	15%	Only path for factoring; government interest ensures funding	Factoring demonstration at RSA-1024 equivalent
Field-theoretic computation	10%	Speculative; fundamental research only	Proof-of-concept field computer for a classically hard problem

4.4 What Is NOT in the Portfolio

Several technologies receive substantial current investment but are absent from our recommended portfolio at near-term and medium-term horizons:

****Universal fault-tolerant gate-model QC (superconducting, trapped ion) as a near/medium-term investment:**** The thermodynamic arithmetic from Phase III shows that these platforms cannot achieve joules-per-solution advantage for any commercially relevant problem within the next decade, except possibly factoring --- and that requires a machine 10⁵× larger than current state-of-the-art. These platforms should be funded as fundamental research, not as commercial ventures. The billions currently flowing into superconducting QC startups represent a capital misallocation that will not produce returns within VC fund lifetimes.

Neuromorphic computing as a general-purpose replacement for GPUs: Neuromorphic processors are specialized for inference and probabilistic computation. They are poor fits for training (which requires backpropagation, not local learning rules) and for general-purpose computation. The appropriate role for neuromorphic is edge inference and specialized sensing, not data center replacement.

****Any technology that cannot state a falsifiable joules-per-solution milestone:**** This is the acid test. If a company cannot state --- in writing, with specific numbers --- the problem class, scale, and joules-per-solution at which their technology will become commercially competitive, their technology is not yet an investment proposition. It is a

research program. Research programs should be funded by research agencies, not by investors seeking financial returns.

5. The Evaluation Framework

For any proposed computational technology --- whether a startup pitch deck, a government grant proposal, or a corporate R&D initiative --- we propose a standardized evaluation rubric:

5.1 The Five Questions

1. **What problem class does it target?** (Optimization, linear algebra, inference, simulation, cryptography, general-purpose)

1. ****What is the physical substrate, and why is it naturally suited to this problem class?**** (Not "what gates does it implement" but "what physics does it exploit")

1. **What is the joules-per-solution at commercially relevant scale?**
(Measured at the wall plug, including all overhead --- cooling, control, error correction, post-processing)

1. ****What is the classical competitor, and at what scale does the crossover occur?**** (Not "conventional CMOS" --- the BEST classical alternative, including specialized hardware: FPGA, ASIC, reversible)

1. **What is the falsifiable milestone with a specific timeframe?**
("We will demonstrate X joules-per-solution advantage on problem Y of commercially relevant scale Z by date W")

5.2 The Red Flags

Any of the following should trigger heightened skepticism:

- Claims of "quantum advantage" without specifying the problem class, the classical competitor, and the joules-per-solution comparison.
- Benchmarks that use random circuit sampling, boson sampling, or other contrived problems with no commercial value.
- Comparisons against unoptimized classical algorithms rather than the best available classical implementation.
- Timelines that have been repeatedly revised outward ("fault-tolerant in 5 years" stated annually since 2015).
- Refusal to engage independent validators who do not have access to proprietary hardware.

5.3 The Green Flags

Conversely, these patterns correlate with genuine progress:

- Publication of SPECIFIC joules-per-solution numbers, not just "quantum volume" or "quantum utility" metrics.
- Engagement with independent validators who publish their own analysis.
- Explicit acknowledgment of the error-correction overhead and its thermodynamic consequences.
- Comparison against specialized classical hardware (FPGAs, ASICs, reversible processors), not just general-purpose CPUs/GPUs.
- Milestones that have been met on or ahead of schedule.

6. Conclusion: The Honest Portfolio

The first four papers in this series have traced an arc from critique to construction:

1. **The Qubit Delusion** identified the epistemic failure: particle ontology projected onto relational reality.
1. **Beyond the Qubit** surveyed alternatives with greater ontological fidelity and commercial manufacturability.
1. **The Physics of Computation** established the honest boundaries imposed by fundamental physical law.
1. **This paper** translates these insights into a concrete investment portfolio and evaluation framework.

The portfolio that emerges is radically different from the current allocation of computational R&D capital. Approximately 45% goes to thermodynamic and analog computing (Ising machines, p-bit networks, analog quantum simulators), 25% to optical/photonic computing, 20% to neuromorphic, and only 5-10% to fault-tolerant quantum --- a near- inversion of the current allocation, where gate-model quantum computing absorbs roughly 60% of advanced computing investment.

This reallocation is not a bet against physics. It is a bet ON physics --- on matching computational problems to the physical substrates that most naturally solve them, rather than forcing all problems into a single, ontologically inappropriate scaffold.

The honest portfolio does not promise exponential speedups or revolutionary new industries within five years. It promises something more valuable: a research program that can actually be falsified, that respects the thermodynamic and information-theoretic limits of computation, and that allocates capital to the places where the physics says it can actually produce returns.

In the language of investment: this is a value portfolio in a field dominated by growth speculation. It will not produce the highest narrative returns. But it may produce the highest actual returns --- measured in joules per solution, not in press releases per quarter.

References

- @Feynman1982: R. P. Feynman, "Simulating Physics with Computers," Int. J. Theor. Phys. 21, 467 (1982).

- @Shen2017: Y. Shen et al., "Deep Learning with Coherent Nanophotonic Circuits," Nat. Photonics 11, 441 (2017).
- @Camsari2017: K. Y. Camsari et al., "Stochastic p-Bits for Invertible Logic," Phys. Rev. X 7, 031014 (2017).
- @Davies2018: M. Davies et al., "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning," IEEE Micro 38, 82 (2018).
- @Mazurenko2017: A. Mazurenko et al., "A Cold-Atom Fermi-Hubbard Antiferromagnet," Nature 545, 462 (2017).
- @Honjo2021: T. Honjo et al., "100,000-Spin Coherent Ising Machine," Sci. Adv. 7, eabh0952 (2021).
- @Gheorghiu2019: C. Gheorghiu et al., "Quantum Resource Estimation for Large Scale Quantum Algorithms," arXiv:1904.01360 (2019).
- @Snider2012: G. L. Snider et al., "Minimum Energy for Computation, Theory vs. Experiment," IEEE Trans. Nanotechnol. 11, 406 (2012).

Changelog

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