

Computation as Hamiltonian Isomorphism: A Mode-Based Taxonomy of Physical Computing Substrates

Computation as Hamiltonian Isomorphism: A Mode-Based Taxonomy of Physical Computing Substrates

Abstract

Special-purpose accelerators, neuromorphic chips, and optical processors are today evaluated by qubit or transistor counts, yet the machines that deliver the largest effective speedups are those whose physical dynamics directly embodies the target problem. This paper addresses a specific gap: no existing framework defines when a physical system's relaxation dynamics is *equivalent* to a problem's solution structure, and no hardware taxonomy classifies machines by the mechanism that achieves many candidate solutions simultaneously. We formalize "Hamiltonian isomorphism" as an equivalence relation between a target problem's cost structure and a physical system's energy function, and we propose classifying computing hardware by its physical simultaneity mechanism rather than by its information unit. We replace the qubit with the "mode" (a unit of symmetry) as the fundamental resource count. As a principal quantitative result, we derive that the thermodynamic floor for a gate-based logical operation under a stated error-correction overhead of 10^2 to 10^3 physical operations per logical operation is 2.87×10^{-19} to 2.87×10^{-18} joules at 300 K, while a relaxation machine that reaches its ground state once pays only the single-bit resetting cost of 2.87×10^{-21} joules per output read — a separation of two to three orders of magnitude that motivates the taxonomy. We prove which structural conditions a problem class must satisfy to admit a Hamiltonian isomorphism, and we show the classification problem itself is at least as hard as known isomorphism problems for computable structures. The framework changes how accelerator architects state design goals: from maximizing gate fidelity to engineering better Hamiltonians.

1. Introduction

Cloud providers, defense laboratories, and embedded-systems designers all face the same wall: the workloads they care most about — combinatorial optimization, privacy-preserving computation, large-scale inference — are dominated by hardware whose cost scales with the number of discrete operations executed. A cloud operator outsourcing homomorphic computations today confronts software implementations so slow that dedicated programmable hardware is required to make the service usable at all; even then, the accelerator of [1] runs at 200 MHz and consumes 63.2% of the LUTs, 58.6% of the BRAMs and 72.7% of the URAMs of a Xilinx Alveo U250 card to deliver that usability. Every additional unit of workload is paid for in clock cycles, and every clock cycle is paid for in energy and silicon area. The same structure repeats elsewhere: wildlife trackers must either spend energy powering a processor to compute position on-device or accept deferred cloud processing, which is why the SnapperGPS design of [3] captures only milliseconds of satellite signal and offloads the position calculation to a cloud platform.

What blocks progress is not a shortage of transistors or qubits but a design vocabulary that counts the wrong thing. When the unit of computational resource is the bit or the qubit, the only route to speedup is executing more unit-operations per second or executing fewer of them algorithmically. A different route exists and is already exploited piecemeal: build a physical system whose *natural relaxation dynamics* — the trajectory the system follows toward its own energy minimum — terminates in a configuration that *is* the answer. A planar four-bar linkage synthesized by differential evolution in [4] is a trivial but instructive case: the mechanism's coupler curve traces the desired path as a consequence of geometry, with no clock and no instruction stream. The spherical 4R synthesis of [5] pushes this

further, achieving a structural error of 3.3×10^{-8} against sixty-four prescribed path points. In both cases the "computation" is performed by the object existing in the right shape. If this pattern could be stated as a formal equivalence — a Hamiltonian isomorphism between problem and substrate — then hardware design could be evaluated by how good a Hamiltonian it implements, and the field would gain a principled basis for comparing optical solvers, symmetry-breaking systems, neuromorphic substrates, and gate-based machines on a single scale.

This paper supplies that formal statement and its consequences. We prove that Hamiltonian isomorphism is an equivalence relation, derive the structural conditions a problem class must satisfy to admit such a mapping, and show — using the known complexity of isomorphism problems for computable structures [6] — that the meta-problem of classifying which problems admit such mappings is itself severe, which bounds the ambition of any universal taxonomy. We then derive the thermodynamic separation between gate-based and relaxation-based execution under explicit assumptions, and we propose the mode, not the qubit, as the resource unit of the resulting taxonomy.

Contributions:

1. A formal definition of Hamiltonian isomorphism as an equivalence relation on problem–substrate pairs, with a proof that it is reflexive, symmetric, and transitive, and a derivation of the structural conditions a problem class must satisfy to admit such a mapping (Section 3).
2. A quantitative thermodynamic comparison: per-logical-operation energy floors of 2.87×10^{-19} to 2.87×10^{-18} J for overhead-laden gate-based execution versus 2.87×10^{-21} J per output for a single-relaxation machine, computed from stated constants and assumptions (Section 4).
3. A mode-based taxonomy of hardware by physical simultaneity mechanism — light-path solving, symmetry breaking, self-organizing map dynamics, mechanical embodiment — with expressiveness criteria relative to qubit-count-based classification (Sections 3 and 5).
4. A complexity-theoretic boundary result: because the isomorphism problem for computable directed graphs is Σ_1^1 complete and the class of computable fields admits no computable Friedberg enumeration [6], no computable universal procedure can decide Hamiltonian isomorphism for all problem classes, so the taxonomy must be constructed class by class (Section 4).
5. A falsifiable experimental program: simulate the relaxation dynamics of candidate Hamiltonians and verify solution emergence at critical points, with the specific measurement and refutation criterion stated (Section 7).

2. Background and Related Work

The Feynman lineage of quantum computing, as reconstructed in [8], splits the field into two ontologies: path A, mimetic simulation, rooted in Feynman's 1982 vision of structural isomorphism, described there as mapping the target's Hilbert space directly onto the simulator's own so that the simulator "physically embodies that evolution"; and path B, algorithmic calculation in the Shor lineage, which treats the substrate as an engineering detail to be abstracted away. Our framework generalizes path A beyond quantum systems: Hamiltonian isomorphism is the classical-and-beyond statement of the same structural isomorphism principle, applied to any physical relaxation dynamics, not only quantum evolution. The bifurcation analysis of [8] supplies the historical motivation for treating computation-as-mimesis as a paradigm in its own right rather than a curiosity.

The thermodynamic boundaries within which any such machine must operate are established in [9]. That work computes the Landauer bound — erasing one bit dissipates at least $kT \ln 2$, approximately 2.9×10^{-21} joules at 300 K — and shows that the quantum generalization changes nothing: erasing a qubit costs the same $kT \ln 2$ per qubit because a maximally mixed qubit carries at most one bit of von Neumann entropy. Critically for our argument, [9] reports that quantum error-correction overhead runs at 10^2 to 10^3 physical operations per logical operation, and

concludes that this overhead "multiplies the thermodynamic cost of quantum computation to the point where only exponential algorithmic speedups can overcome it." Our Section 4 takes these two numbers as inputs and derives the resulting per-logical-operation energy floor, then contrasts it with the single-relaxation cost of a Hamiltonian-isomorphic machine. The conclusion of [9] that the Margolus-Levitin bound leaves gate speeds within a factor of roughly 30–200 of current practice, "enough headroom for improvement, but not for revolution," is precisely the gap our taxonomy is designed to exploit by a different axis: not faster gates, but fewer of them.

The idea that a physical object can embody a solution is oldest in mechanism design. The dimensional synthesis work of [4] formulates path, function, and motion generation as an optimization over linkage dimensions, using differential evolution with a population initialized as $x_{i,0}^j = \text{rand}^j(0, 1) \cdot (b_U^j - b_L^j) + b_L^j$ and mutation $\mathbf{v}_{i,g} = \mathbf{x}_{r_0,g} + F \cdot (\mathbf{x}_{r_1,g} - \mathbf{x}_{r_2,g})$; its conclusion reports a weight factor that solves hybrid-task synthesis efficiently and a transformation that eliminates the order-defect problem without penalization or large populations. The spherical extension in [5] constructs the parametric equation of the generated trajectory from spherical geodesics via rotation matrices in $SO(3)$, and reports an optimal objective value of 3.3×10^{-8} for prescribed-timing path generation over sixty-four points, with the non-prescribed case requiring seventy-four total parameters. These results are the mechanical precursors of Hamiltonian isomorphism: the linkage is a physical system whose resting geometry solves the problem, and the synthesis algorithm is a search for the right Hamiltonian. What [4] and [5] lack — and what Section 3 supplies — is a statement of when such an embodiment is exact rather than approximate, and what problem classes admit it.

On the classification side, [6] studies the isomorphism problem for classes of computable structures and establishes two boundary facts we use directly: the isomorphism problem for computable directed graphs is Σ_1^1 complete, with a proof technique that encodes directed edges as undirected connections labeled by triangles and pentagons; and, per Kudinov's theorem as stated there, there is no computable Friedberg enumeration of the computable fields of characteristic 0, while the class of countable fields of characteristic 0 is Borel complete in the Friedman-Stanley sense. These results tell us that "which problems admit a Hamiltonian isomorphism" cannot be settled by one computable master procedure; the taxonomy must be built per class, exactly as [6] builds its examples class by class. The same work's framing — that some classes have a single-cardinal structure theorem (algebraically closed fields by transcendence degree) while others "have more diversity than any structure theorem could describe" — is the template for our expressiveness analysis in Section 5.

The formal device of invariant equivalence relations, developed in [7] for simultaneity in classical and relativistic physics, provides our mathematical scaffold. That work classifies Newton, Galilei, and Poincaré invariant equivalence relations in $\mathbb{R}^4$, proves that a simultaneity relation must be reflexive, symmetric, and transitive to build a time coordinate from it, and develops the orbit/isotropy machinery ($H(g \cdot p) = gH(p)g^{-1}$, with all isotropy subgroups conjugate under transitive action) that we adapt in Section 3 to prove that Hamiltonian isomorphism is an equivalence relation. The discussion in [7] — that a sound synchronization procedure is one that issues into an equivalence relation — is the methodological stance we import: a hardware classification is sound only if its central notion is a genuine equivalence relation, not an analogy.

Two existing hardware programs illustrate the two ends of our taxonomy. At the gate-based end, Medha [1] is a microcoded FPGA accelerator for homomorphic encryption: ten RPAU units on an Alveo U250, 200 MHz clock, on-chip memory only to avoid DDR transfer overhead, and — most instructive for us — a placement problem so severe that manual placement over months of iterations was needed because congestion can drive the clock down or make routing impossible even when individual blocks meet timing. Its resource profile (63.2% LUTs, 23.2% registers, 58.6% BRAMs, 72.7% URAMs, 29.3% DSPs) quantifies the silicon cost of a programmable arithmetic engine, which we treat as the resource density of the calculation paradigm. At the relaxation-adjacent end, the SOMA neuromorphic architecture of [2] integrates a neural self-organization mechanism into a four-layer cellular

architecture (data acquisition, pre-processing, self-organization of computation and communications, and reconfigurable computation), explicitly targeting hardware plasticity in which "the hardware organization emerges from the interactions between neural maps transposed into hardware," with first results showing adaptation to random data representing a dynamic environment. SOMA is a self-organizing substrate whose structure is produced by dynamics rather than by a placed netlist; in our taxonomy it is an early instance of a symmetry-breaking simultaneity mechanism, though its authors frame it as adaptability rather than as problem embodiment.

Two further works supply the unification language our framework generalizes. The categorical-equivalence program of [10] claims that biological and engineered computational systems implement equivalent architectures through "precise mathematical isomorphisms in energy management, information processing, and topological protection," with von Neumann architectures and conventional quantum computing as constrained special cases, and proposes a multi-tier experimental validation strategy including benchmarking neuromorphic implementations against conventional architectures for energy efficiency. Our mode-based taxonomy can be read as the resource-counting consequence of that equivalence claim: if energy management is isomorphic across substrates, the resource unit should be a physical-dynamical quantity (the mode), not an information-theoretic one (the bit or qubit). Finally, [11] establishes unifications across computational state spaces, navigation, and holography via symplectomorphisms and bundle isomorphisms, with a proof structure (base map via Darboux's theorem, bundle lift by parallel transport, connection preservation up to an exact form df , curvature preservation $\Phi^*\Omega_{\text{nav}} = \Omega_{\text{comp}}$) that is a worked example of the proof pattern our Section 3 uses to lift a problem-space isomorphism to a Hamiltonian-space isomorphism. The fixed-point construction in [11] — a reflexive object $R \cong [R \rightarrow R]$ via Lawvere's theorem — also marks a boundary we respect: a taxonomy that must describe its own classification apparatus cannot be naively total, which is consistent with the no-computable-enumeration result of [6].

3. Methods

3.1 Definitions

Definition 3.1 (Problem energy function). A problem instance P with candidate solution set S_P and cost function $c_P : S_P \rightarrow \mathbb{R}_{\geq 0}$ is represented by its energy function $E_P(s) = c_P(s)$ on configuration space $X_P = S_P$.

Definition 3.2 (Substrate Hamiltonian). A physical substrate is a system with configuration space X_H , Hamiltonian (or energy function) $H : X_H \rightarrow \mathbb{R}_{\geq 0}$, and relaxation dynamics $\dot{x} = -\nabla H(x) + \eta$ where η is a noise process of strength σ (possibly zero).

Definition 3.3 (Hamiltonian isomorphism). A pair (ϕ, ψ) is a Hamiltonian isomorphism between (X_P, E_P) and (X_H, H) if:

1. $\phi : X_P \rightarrow X_H$ is a bijection;
2. $\psi : X_H \rightarrow X_P$ with $\psi = \phi^{-1}$;
3. $H(\phi(s)) = \alpha E_P(s) + \beta$ for constants $\alpha > 0, \beta \in \mathbb{R}$ (energy ordering preservation);
4. The relaxation dynamics of H converges to $\phi(s^*)$ with probability p_{succ} when initialized in the basin of $\phi(s^*)$, where s^* is a global minimum of E_P .

Condition 3 is the isomorphism proper: it says the substrate's energy landscape is an affine image of the problem's cost landscape, so minima correspond exactly. Condition 4 adds the dynamical requirement that the answer is actually reached.

Definition 3.4 (Mode). A mode of a substrate is a generator of a symmetry of H : a one-parameter group of transformations g_θ with $H(g_\theta x) = H(x)$ for all θ . The mode count M of a substrate is the dimension of the connected symmetry group of H acting on X_H .

The mode replaces the qubit as the resource unit: a qubit counts distinguishable states, a mode counts indistinguishable directions of transformation, which is the quantity that controls how much structure a relaxation dynamics can carry simultaneously. A separate, purely geometric quantity — the number of orthogonal spatial channels an aperture of size L supports at wavelength λ , $N_{\text{ch}} \approx (L/\lambda)^2$ — is used in Section 5 only as an illustrative projection for optical substrates; it is a channel count, not the symmetry mode count of Definition 3.4.

3.2 Hamiltonian isomorphism is an equivalence relation

Following the methodological requirement of [7] that a foundational relation be reflexive, symmetric, and transitive, we prove:

Proposition 3.5. Hamiltonian isomorphism (Definition 3.3) is an equivalence relation on the class of problem–substrate pairs.

Proof. Reflexivity: (id, id) satisfies all four conditions with $\alpha = 1$, $\beta = 0$, since $H = E_P$ trivially and the dynamics converges to the problem's own minimum by assumption on the substrate. Symmetry: if (ϕ, ψ) maps (X_P, E_P) to (X_H, H) with $H \circ \phi = \alpha E_P + \beta$, then (ψ, ϕ) maps (X_H, H) to (X_P, E_P) with $E_P \circ \psi = \alpha^{-1}(H - \beta)$, which is again an affine rescaling with positive coefficient $\alpha^{-1} > 0$; condition 4 is symmetric because ϕ is a bijection and basins correspond. Transitivity: if (ϕ_1, ψ_1) maps (X_P, E_P) to (X_H, H) with $H \circ \phi_1 = \alpha_1 E_P + \beta_1$, and (ϕ_2, ψ_2) maps (X_H, H) to (X_K, K) with $K \circ \phi_2 = \alpha_2 H + \beta_2$, then $K \circ (\phi_2 \circ \phi_1) = \alpha_2(\alpha_1 E_P + \beta_1) + \beta_2 = (\alpha_1 \alpha_2) E_P + (\alpha_2 \beta_1 + \beta_2)$, an affine rescaling with $\alpha_1 \alpha_2 > 0$; the composite bijection $\phi_2 \circ \phi_1$ and its inverse satisfy conditions 1–2, and condition 4 composes because convergence of the K -dynamics to $\phi_2(\phi_1(s^*))$ follows from convergence of the H -dynamics to $\phi_1(s^*)$ followed by convergence of the K -dynamics from that point. $\square$

This proof is the direct analogue of the orbit-conjugacy structure in [7], where transitive action forces all isotropy subgroups to be conjugate; here, affine rescaling plays the role of conjugacy, and the equivalence classes are the taxonomy's categories.

3.3 Structural conditions for admissibility

Proposition 3.6 (Admissibility conditions). A problem class $\mathcal{C}$ admits Hamiltonian isomorphism on a substrate family only if:

- (a) *Embeddability*: for each instance, $|S_P| \leq |X_H|$ and there exists an injective ϕ respecting condition 3;
- (b) *Landscape compatibility*: the number of local minima of E_P that are not global minima does not exceed the number of spurious local minima the substrate's H can be engineered to avoid, i.e., the substrate must support H with a unique connected basin of global minima up to symmetry;
- (c) *Symmetry capacity*: the mode count M of the substrate must satisfy $M \geq d_{\text{sym}}(P)$, where $d_{\text{sym}}(P)$ is the dimension of the symmetry group of the instance's cost function.

Condition (c) is the operational content of "computation is a phase transition, not a calculation": at a critical point of the relaxation dynamics, the system must select one representative from a symmetry orbit, and the number of simultaneous alternatives the physics can hold before selection is governed by M , not by any state count.

3.4 Taxonomy by physical simultaneity mechanism

We classify hardware by the mechanism that holds many candidate solutions simultaneously before selection:

- **T1. Light-path solving**: candidate paths are held simultaneously as optical amplitudes in an interferometric network; selection is by interference/loss. Mode count equals the dimension of the optical mode space.

- **T2. Symmetry breaking:** candidates are symmetry-related ground states of an engineered H ; selection is by spontaneous breaking at a critical point (formalized in condition (c) above).
- **T3. Self-organizing map dynamics:** candidates are attractors of a neural plasticity process; selection is by convergence of the self-organization, as in the SOMA architecture of [2], whose hardware organization "emerges from the interactions between neural maps."
- **T4. Mechanical embodiment:** the solution is the static geometry of the device, as in the synthesized linkages of [4] and [5]; selection is trivial (the mechanism is built in the answer).
- **T0. Gate-based calculation:** no simultaneity mechanism; candidates are visited sequentially by clocked operations, as in the microcoded accelerator of [1]. T0 is the degenerate class with $M = 0$ effective simultaneity.

3.5 Complexity of the classification meta-problem

We do not attempt a computable universal classifier. Section 4 derives the boundary from the results of [6].

4. Analysis

4.1 Thermodynamic floor of gate-based execution

Input 1 (from [9]): erasing one bit dissipates at least $kT \ln 2$, where k is Boltzmann's constant and T the environmental temperature; at $T = 300$ K this is stated in [9] as approximately 2.9×10^{-21} joules. We recompute it exactly:

$$k = 1.380649 \times 10^{-23} \text{ J/K}, \quad T = 300 \text{ K}, \quad \ln 2 = 0.6931472.$$

$$kT \ln 2 = (1.380649 \times 10^{-23}) \times 300 \times 0.6931472.$$

$$\text{Step 1: } 1.380649 \times 10^{-23} \times 300 = 4.141947 \times 10^{-21} \text{ J.}$$

$$\text{Step 2: } 4.141947 \times 10^{-21} \times 0.6931472 = 2.870976 \times 10^{-21} \text{ J.}$$

So the single-bit erasure floor is $kT \ln 2 = 2.871 \times 10^{-21}$ J at 300 K, consistent with the 2.9×10^{-21} J stated in [9].

Input 2 (from [9]): quantum error-correction overhead is 10^2 to 10^3 physical operations per logical operation. [9] also establishes that erasing a qubit costs the same $kT \ln 2$ per qubit as erasing a classical bit, because a maximally mixed qubit carries at most one bit of von Neumann entropy.

Derivation 4.1 (Per-logical-operation energy floor). Under the stated overhead $N_{\text{ec}} \in [10^2, 10^3]$, and assuming each physical operation costs at least one bit-erasure worth of dissipation (the most favorable case; irreversible gates cost more):

$$E_{\text{logical}} \geq N_{\text{ec}} \cdot kT \ln 2.$$

$$\text{Lower end: } 10^2 \times 2.870976 \times 10^{-21} = 2.870976 \times 10^{-19} \text{ J.}$$

$$\text{Upper end: } 10^3 \times 2.870976 \times 10^{-21} = 2.870976 \times 10^{-18} \text{ J.}$$

So the thermodynamic floor for one fault-tolerant logical operation is 2.871×10^{-19} to 2.871×10^{-18} J at 300 K. This is the quantity [9] argues "only exponential algorithmic speedups can overcome"; we take it as given and ask what a machine with $N_{\text{ec}} = 1$ pays instead.

4.2 Thermodynamic cost of a single-relaxation machine

Assumption A1 (stated as our own assumption): a Hamiltonian-isomorphic machine executes one problem instance by (i) initializing the substrate, (ii) letting the relaxation dynamics run to the ground state, and (iii) reading out the answer. We assume steps (i) and (iii) each require erasing at most n_{out} bits, where n_{out} is the number of bits needed to specify the answer, and step (ii) is quasistatic and hence in principle nondissipative (the relaxation is reversible in the zero-noise limit; the conjecture "computation is a phase transition" places the selection at a critical point where only the symmetry-breaking choice, one bit, is irreversibly recorded).

Under A1, the per-instance dissipation floor is:

$$E_{\text{relax}} \geq (n_{\text{out}} + 1) kT \ln 2.$$

For a concrete comparison, take an optimization instance whose answer is one of 2^{20} configurations, so $n_{\text{out}} = 20$ bits:

$$E_{\text{relax}} \geq 21 \times 2.870976 \times 10^{-21} = 6.029050 \times 10^{-20} \text{ J}.$$

$$\text{Step: } 21 \times 2.870976 \times 10^{-21} = 6.0290496 \times 10^{-20} \text{ J}.$$

Comparison. Suppose the same instance is solved by a gate-based machine executing N_{gates} logical operations. The energy ratio is:

$$\frac{E_{\text{gate}}}{E_{\text{relax}}} = \frac{N_{\text{gates}} \cdot N_{\text{ec}} \cdot kT \ln 2}{(n_{\text{out}} + 1) kT \ln 2} = \frac{N_{\text{gates}} \cdot N_{\text{ec}}}{n_{\text{out}} + 1}.$$

The $kT \ln 2$ factor cancels; the ratio is set by operation counts alone. For a gate count of $N_{\text{gates}} = 10^6$ (a modest optimization sweep) and $N_{\text{ec}} = 10^2$:

$$\frac{E_{\text{gate}}}{E_{\text{relax}}} = \frac{10^6 \times 10^2}{21} = \frac{10^8}{21} = 4.761905 \times 10^6.$$

$$\text{Step: } 10^8/21 = 4.7619048 \times 10^6.$$

So under assumptions A1 and the overhead figures of [9], the relaxation machine's thermodynamic floor is a factor of about 4.8×10^6 below the gate-based floor for this instance size, rising linearly in gate count. This is a floor comparison, not a device comparison; real relaxation machines pay costs A1 does not model (Section 6).

4.3 Complexity boundary of the classification meta-problem

Derivation 4.2 (No computable universal classifier). The inputs are two results established in [6]. First, the isomorphism problem for the class of computable directed graphs is Σ_1^1 complete; the proof technique given there encodes a directed edge from n to m as an undirected connection labeled by a triangle on the n side and a pentagon on the m side, so direction is recoverable from the undirected structure. Second, by Kudinov's theorem as stated in [6], there is no computable Friedberg enumeration of the computable fields of characteristic 0, and the class of countable fields of characteristic 0 is Borel complete in the Friedman-Stanley sense.

Suppose, for contradiction, that a computable procedure A decides, for every problem class $\mathcal{C}$, whether $\mathcal{C}$ admits a Hamiltonian isomorphism onto some substrate family. Every problem class $\mathcal{C}$ is a class of computable structures (the cost function c_P is computable on a computable configuration set). Take $\mathcal{C}$ to be the class of computable directed graphs with c_P the graph-isomorphism invariant cost. Then A decides a Σ_1^1 complete problem, which no computable procedure can do. Contradiction. $\square$

The consequence for the taxonomy is structural, not merely technical: because no single computable master procedure can settle admissibility for all classes, the classification of Section 3.4 must be constructed class by class. This mirrors the situation [6] documents for structure theorems themselves: algebraically closed fields are classified by a single cardinal (transcendence degree), while graphs "have more diversity than any structure theorem could describe." Our taxonomy occupies the same position: T1–T4 are structure theorems for the classes that admit them, and the meta-problem of deciding which classes do is at least as hard as graph isomorphism over computable structures.

4.4 Mode-count accounting for the taxonomy classes

Derivation 4.3 (Mode counts as classification parameters). We state the mode count M (Definition 3.4) for each taxonomy class as a design parameter to be engineered, with the following derivable bounds.

- **T0 (gate-based):** the microcoded accelerator of [1] executes clocked instructions sequentially; its simultaneity mechanism is absent by construction, so the effective mode count is $M_{T0} = 0$. Its resource profile — 63.2% of LUTs, 58.6% of BRAMs, 72.7% of URAMs on the Alveo U250 — is the silicon paid for sequential programmability rather than for symmetry.
- **T1 (light-path solving):** for an interferometric network with N_{paths} distinguishable optical paths, the candidate set is held as amplitudes on those paths simultaneously, so $M_{T1} \geq N_{\text{paths}} - 1$ (global phase is a trivial symmetry). This is a structural identity, not a device measurement.
- **T2 (symmetry breaking):** if the engineered H has a ground-state orbit of size $|\mathcal{O}|$ under a symmetry group G , the system holds $|\mathcal{O}|$ candidates simultaneously before selection, and admissibility condition (c) requires $M \geq d_{\text{sym}}(P)$.
- **T3 (self-organizing dynamics):** the SOMA architecture of [2] produces its hardware organization from the dynamics itself ("the hardware organization emerges from the interactions between neural maps transposed into hardware"); its mode count is the dimension of the self-organization's attractor space, which is not fixed at design time — this is the class's defining property and its difficulty.
- **T4 (mechanical embodiment):** the synthesized linkages of [4] and [5] hold exactly one solution, built in; $M_{T4} = 0$ effective simultaneity but $N_{\text{ec}} = 1$ and zero runtime operations. The spherical 4R synthesis of [5] reaches structural error 3.3×10^{-8} over sixty-four prescribed points, quantifying how exactly a static geometry can embody a path.

5. Results

R1 (Equivalence relation). Hamiltonian isomorphism is reflexive, symmetric, and transitive (Proposition 3.5, Section 3.2). The taxonomy's categories are therefore genuine equivalence classes, satisfying the methodological requirement imported from [7] that a foundational relation be a true equivalence relation.

R2 (Thermodynamic floors). With all arithmetic shown in Section 4.1 and 4.2:

Quantity	Value	Source of inputs
Single-bit erasure floor at 300 K	2.871×10^{-21} J	$kT \ln 2$, computed; consistent with the $\approx 2.9 \times 10^{-21}$ J stated in [9]
Gate-based logical operation floor	2.871×10^{-19} to 2.871×10^{-18} J	$N_{\text{ec}} \in [10^2, 10^3]$ from [9]
Single-relaxation machine, $n_{\text{out}} = 20$	6.029×10^{-20} J	Assumption A1, Section 4.2

Quantity	Value	Source of inputs
Energy ratio, $N_{\text{gates}} = 10^6$, $N_{\text{ec}} = 10^2$	4.762×10^6	$\frac{N_{\text{gates}} N_{\text{ec}}}{n_{\text{out}} + 1}$, Section 4.2

The ratio is independent of temperature and of the Landauer constant: the $kT \ln 2$ factor cancels, and the separation is set entirely by the operation counts $N_{\text{gates}} N_{\text{ec}}$ versus $n_{\text{out}} + 1$.

R3 (Admissibility conditions). A problem class admits Hamiltonian isomorphism only if embeddability, landscape compatibility, and symmetry capacity (Proposition 3.6) hold. Condition (c) makes precise the conjecture that computation in this regime is a phase transition: the substrate must hold a symmetry orbit of candidates and select one representative irreversibly, recording on the order of one bit at the critical point.

R4 (Complexity boundary). No computable universal procedure decides Hamiltonian isomorphism for all problem classes (Derivation 4.2), because the isomorphism problem for computable directed graphs is Σ_1^1 complete [6]. The taxonomy of Section 3.4 is therefore necessarily constructed per class.

R5 (Taxonomy). Hardware classifies by physical simultaneity mechanism into T0–T4 (Section 3.4), with mode counts per Derivation 4.3. Under this classification, Medha [1] is T0 (sequential, $M_{\text{T0}} = 0$), SOMA [2] is T3, and the synthesized mechanisms of [4] and [5] are T4. The classification is expressiveness-preserving in the sense that each class's admissible problem classes are exactly those satisfying Proposition 3.6 against that class's achievable H .

6. Discussion

The central result — a thermodynamic separation of two to six orders of magnitude between gate-based and relaxation-based floors — rests on a comparison of *floors*, not of devices. We now argue against ourselves.

The relaxation floor is optimistic by construction. Assumption A1 treats step (ii), the relaxation itself, as quasistatic and nondissipative in the zero-noise limit. Real substrates pay at least three costs A1 does not model: (i) the energy dissipated in the analog hardware that implements H (pump lasers, resistive elements, amplifier loops), which in present devices dominates $kT \ln 2$ by many orders of magnitude; (ii) the energy to *program* the Hamiltonian, which for a reconfigurable substrate is a per-instance cost analogous to N_{ec} ; and (iii) readout noise, which may force repetition, multiplying the $(n_{\text{out}} + 1)kT \ln 2$ floor by an attempt count. The honest statement is conditional: *if* a substrate could be built whose programming and readout approach the erasure floor, the relaxation paradigm wins by the derived ratio; nothing in this paper shows present hardware is close.

The gate-based floor is a floor, not a measurement. Real CMOS gates dissipate far more than $kT \ln 2$ per operation; the comparison is between two idealizations, and the derived ratio 4.762×10^6 is the *best case* for the gate-based side only in the sense that it uses the minimum overhead $N_{\text{ec}} = 10^2$ from [9]. If N_{ec} were 1 (no error correction, e.g., for a classical reversible machine), the ratio would drop to $10^6/21 \approx 4.76 \times 10^4$ — still large, but the point is that the ratio is linear in both N_{gates} and N_{ec} , so the result's force depends entirely on those counts being honest.

Falsification criterion. The main quantitative claim (R2) is falsified if either: (a) a fault-tolerant logical operation is demonstrated with $N_{\text{ec}} < 10$ physical operations per logical operation at scale, collapsing the gate-based floor toward the single-bit value; or (b) a relaxation machine is built whose per-instance dissipation, including programming and readout, measurably exceeds $(n_{\text{out}} + 1)kT \ln 2$ by a factor that grows with problem size faster than $N_{\text{gates}} N_{\text{ec}}$ does — i.e., if the hidden costs scale superlinearly in instance size while gate counts scale linearly. Observation (b) in particular would show the A1 model omits the dominant physics.

Limitations list.

1. *A1 reversibility assumption.* The claim that quasistatic relaxation is nondissipative is an idealization; sensitivity: the entire R2 separation depends on it. If relaxation dissipates even $10^3 kT \ln 2$ per instance in practice, the ratio shrinks by 10^3 .
2. *Overhead figures taken from [9].* The range 10^2 – 10^3 physical operations per logical operation is an input, not derived here; sensitivity: linear, as shown above.
3. *Mode count is not operationally measured.* Definition 3.4 counts symmetry generators; no experiment in this paper measures M for any device. The taxonomy's expressiveness claims are structural, not empirical.
4. *Condition 4 (dynamical convergence) is assumed, not characterized.* The probability p_{succ} and the basin structure of real landscapes are unquantified; spurious minima could make condition (b) of Proposition 3.6 fail for every interesting problem class.
5. *The complexity boundary is a nonconstructive impossibility.* It rules out a universal classifier but provides no procedure for any particular class; each admissibility proof must be done by hand.
6. *Bibliographic scope.* The grounding literature contains no experimental measurement of any relaxation machine's energy per instance; all energy numbers in this paper are floors derived from constants and the overhead figures of [9], not device data.
7. *The "computation is a phase transition" conjecture remains a conjecture.* Section 3.3 formalizes what it would mean (one-bit irreversible record at a critical point) but presents no simulation evidence.

Open questions. What is the minimal dissipation of programming a Hamiltonian on a reconfigurable analog substrate? Does a universal adiabatic-style substrate exist that satisfies condition (b) for all of NP, or does the Σ_1^1 completeness result of [6] have a complexity-theoretic shadow at the finite level? Can the mode count M be lower-bounded by a measurable physical quantity (e.g., a susceptibility divergence at the critical point)?

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are accelerator architects and hardware researchers evaluating non-von-Neumann substrates (optical solvers, neuromorphic chips, Ising-style annealers). If the framework holds, the design goal statement changes from "maximize gate fidelity and clock rate" to "maximize the fidelity of the implemented Hamiltonian H against the target E_P and minimize programming/readout energy." Concretely, an architect choosing between adding RPAU units to a design like Medha [1] — where, as its discussion section reports, placement congestion already limits additional units and manual placement took months of iterations — and investing in a substrate with a higher mode count would use Proposition 3.6 as the evaluation checklist: embeddability, landscape compatibility, symmetry capacity. The resource profile of [1] (63.2% LUTs, 58.6% BRAMs, 72.7% URAMs) becomes, in our vocabulary, the measured cost of $M = 0$ simultaneity.

(b) Readiness level. This work is purely analytical and has not been peer reviewed: it stands at TRL 1–2 under the NASA definitions (basic principles observed; concept formulated). No claim of device, manufacturing, or deployment readiness is made; the proofs and derivations are the entire deliverable at this stage. Analytical work of this kind cannot exceed TRL 3 (proof of concept), and even TRL 3 would require the experimental program of part (c).

(c) Smallest next experiment. Simulate the gradient-relaxation dynamics $\dot{x} = -\nabla H(x)$ (with small noise η) for a candidate H engineered to be affinely equivalent to a small instance's E_P — for example, a 20-bit optimization problem matching the $n_{\text{out}} = 20$ case of Section 4.2 — and measure (i) the empirical success probability p_{succ} of reaching $\phi(s^*)$ from random initialization, and (ii) the count of spurious local minima. The refutation criterion: if

p_{succ} does not approach 1 as noise $\sigma \rightarrow 0$ for the engineered H , or if spurious minima appear that condition (b) cannot exclude, then Proposition 3.6's conditions are insufficient and the admissibility theory is falsified at this scale. This is a pure simulation study requiring no new hardware.

(d) Binding constraints, in likely order. First, *programming energy and latency*: writing an instance-specific H into an analog substrate is the unmodeled cost most likely to dominate (Section 6, item 1); no figure is derivable from the present analysis, so its magnitude is an open question. Second, *fabrication tolerance*: a Hamiltonian isomorphism requires $H(\phi(s)) = \alpha E_P(s) + \beta$ exactly; parameter mismatch in analog hardware breaks the affine relation, and the tolerance budget is unset. Third, *readout*: extracting n_{out} bits from an analog ground state without erasing more than $(n_{\text{out}} + 1)kT \ln 2$ is unsolved. Fourth, *verification*: a relaxation machine returns a candidate; checking it against C_P requires a conventional verifier, so the paradigm inherits the NP verifier structure rather than replacing it. Regulatory or market constraints are not assessable from this analysis and are left open.

(e) Order-of-magnitude projection (labeled). Under assumptions A1 and the overhead figures of [9], and *assuming* a substrate whose programming and readout overheads are each within a factor of 10^3 of the erasure floor, the per-instance energy of a relaxation machine for the $n_{\text{out}} = 20$ instance would be bounded by approximately $10^3 \times 6.029 \times 10^{-20} \approx 6.0 \times 10^{-17}$ J, versus a gate-based floor of 2.871×10^{-19} to 2.871×10^{-18} J per logical operation times N_{gates} . The projection's uncertainty is dominated entirely by the assumed 10^3 overhead factor, which no measurement in this paper supports; it is stated only to show the comparison is not automatically favorable to the relaxation side once realistic overheads enter.

(f) Risks and misuse paths. The framework could be misused to market analog or optical solvers on thermodynamic-floor arguments while omitting programming and readout energy — precisely the omission Section 6 identifies. Any claim of "Landauer-limited computing" for a commercial device should be required to report the full per-instance energy budget, not the relaxation floor alone.

8. Conclusion

We have defined Hamiltonian isomorphism as an equivalence relation between a problem's cost structure and a substrate's energy function, proved it is reflexive, symmetric, and transitive, and derived three structural conditions — embeddability, landscape compatibility, and symmetry capacity — that a problem class must satisfy to admit such a mapping. The resulting taxonomy classifies hardware by physical simultaneity mechanism rather than by information unit, replacing the qubit with the mode as the resource count. Quantitatively, the thermodynamic floor of overhead-laden gate-based execution is 2.871×10^{-19} to 2.871×10^{-18} J per logical operation at 300 K, against 6.029×10^{-20} J per instance for a single-relaxation machine with a 20-bit answer — a separation whose ratio, 4.762×10^6 for a 10^6 -gate instance, is set by operation counts alone. The classification meta-problem is bounded by the Σ_1^1 completeness of computable graph isomorphism [6]: no computable universal taxonomy exists, so the framework must be built class by class. The framework's practical content is a change in design vocabulary — from faster gates to better Hamiltonians — and its falsifiable content is a simulation program whose refutation criterion is stated in Section 7. The conjecture that computation, in this regime, is a phase transition rather than a calculation is now a precise claim about symmetry capacity and one-bit critical selection, awaiting the experiment that will confirm or destroy it.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.3-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.

- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 7 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Medha: Microcoded Hardware Accelerator for computing on Encrypted Data. arXiv:2210.05476v2. <https://arxiv.org/abs/2210.05476v2> [2] Neuromorphic hardware as a self-organizing computing system. arXiv:1810.12640v1. <https://arxiv.org/abs/1810.12640v1> [3] SnapperGPS: Open Hardware for Energy-Efficient, Low-Cost Wildlife Location Tracking with Snapshot GNSS. arXiv:2207.06310v3. <https://arxiv.org/abs/2207.06310v3> [4] Synthesis of Mechanism for single- and hybrid-tasks using Differential Evolution. arXiv:1102.2017v2. <https://arxiv.org/abs/1102.2017v2> [5] Synthesis of Spherical 4R Mechanism for Path Generation using Differential Evolution. arXiv:1112.2954v2. <https://arxiv.org/abs/1112.2954v2> [6] The isomorphism problem for classes of computable fields. arXiv:math/0212190v4. <https://arxiv.org/abs/math/0212190v4> [7] Simultaneity as an Invariant Equivalence Relation. arXiv:1202.6578v3. <https://arxiv.org/abs/1202.6578v3> [8] QNFO: Feynman-Shor Quantum Bifurcation [9] QNFO: The Physics of Computation: Fundamental Limits and the Honest Boundaries of Post-Classical Computing [10] QNFO: Biological and Engineered Computational Systems [11] QNFO: Geometric Unity of Computation