V1.0.0

A Noether-Type Correspondence for Physical Computing: Continuous Symmetry as a Conserved Computational Resource in Hybrid Intelligence

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Physical computing substrates—resonator networks, cyclic mechanical structures, and engineered Hamiltonian systems—promise order-of-magnitude energy savings over digital inference, but no existing principle tells a designer which physical structure to build so that its symmetry does computational work. This paper addresses that gap. We formalize a correspondence between a continuous $SO(2)/U(1)$ symmetry of a physical computing substrate and a conserved computational quantity: a sector label (harmonic index) that the substrate dynamics cannot mix, analogous to Noether's theorem relating continuous symmetries to conservation laws. Exploiting the conservation law, we derive a block-diagonalization cost model in which an N_s-fold cyclic substrate reduces a dense spectral problem of size $M = N_s m$ from $c\,M^3$ to $N_s\,c\,m^3$ operations, a speedup of N_s^2 (computed: $64\times$ for $N_s = 8$, $m = 100$), and a parameter-sharing reduction of factor m for cyclic layers (computed: $65{,}536 \to 256$ parameters at $m = 256$). Against the Landauer bit energy $k_B T \ln 2 = 2.87 \times 10^{-21}\,\mathrm{J}$ at $T = 300\,\mathrm{K}$, we frame the gains as bits-per-joule headroom under the accounting conventions of thermodynamic epiplexity. The result gives hybrid-system designers a concrete selection rule—maximize exploitable continuous symmetry—and defines the smallest falsifying test: a symmetry-constrained autoencoder whose parameter count and reconstruction fidelity are compared against an unconstrained baseline.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). A Noether-Type Correspondence for Physical Computing: Continuous Symmetry as a Conserved Computational Resource in Hybrid Intelligence. QNFO Preprint v1.0.0.

1. Introduction

Artificial intelligence at the network edge faces a hard resource wall. The survey of edge intelligence by [1] documents that AI services must run on billions of mobile and IoT devices in resource-constrained environments, and that even configuring an edge inference system becomes "a high dimensional configuration problem that is required to determine a large number of performance-critical configuration parameters in real time," with video-analytics parameters including frame rate, resolution, model selection and model early-exit. Every joule and every parameter saved at the edge translates directly into latency, battery life, and privacy-preserving on-device capability, since [1] also identifies data privacy as the key challenge that pushes training and inference away from centralized clouds. Meanwhile, [4] reports that large-scale training runs consume megawatt-hours of electricity while the human brain achieves high-level cognition at roughly $20\,\mathrm{W}$, an orders-of-magnitude efficiency gap that motivates defining intelligence by bits-per-joule rather than benchmark scores alone.

Physical computing—using an engineered Hamiltonian's natural dynamics as the computational step, as formalized in the Inverted Church-Turing-Deutsch framework of [9]—is one proposed escape from that wall. But a designer choosing a substrate today has no first-principles rule for what geometry the substrate should have. The choice between a generic coupled network and a structured one (a ring of resonators, a cyclic crystal, a rotational metamaterial) is cur-

rently made by intuition and per-domain craft. If a general principle existed—stating which symmetries of a physical substrate yield provable computational savings and what quantity those symmetries conserve—then substrate selection, co-design of the physical device with its learning architecture, and the hybrid orchestration layer around it (the dataflow question raised by [2] and [5]) would all acquire a shared design variable.

This paper proposes and develops such a principle. Classical mechanics already supplies the template: when a Lagrangian has a cyclic coordinate, Routh reduction removes that degree of freedom, and [3] extends this reduction to hybrid systems with continuous and discrete dynamics, showing that a cyclic Lagrangian and a cyclic constraint function jointly allow reduction of a hybrid Lagrangian system via a Routhian $R_c^\mu : TP \to \mathbb{R}$. We transport this structure into computing: the substrate plays the role of the mechanical system, the computational state space plays the role of the configuration manifold, and the conserved momentum plays the role of a conserved sector label that the computation cannot leave. The claim is not metaphorical but structural: for any substrate whose update rule commutes with an $SO(2)$ action, the harmonic index is exactly conserved, block decomposition is exact, and the cost reductions follow from linear algebra alone.

Where two cost conventions were available for the worked example—one counting dense matrix–vector updates per step, one counting dense spectral solves—we take the dense spectral-solve convention under which the speedup is N_s^2; the alternative per-step convention gives the same qualitative conclusion with a different exponent, and the choice is stated here so that every number below is unambiguous.

The contributions are:

- A Noether-type correspondence for computation (derived in Section 3): a continuous symmetry G of a substrate's update map implies a conserved computational quantity—the irreducible-representation label of the state—and the update dynamics decompose into invariant blocks indexed by that label.
- An exact cost model for cyclic substrates (derived in Section 4): for N_s sectors of size m, the spectral-problem cost falls from $c\,M^3$ to $N_s\,c\,m^3$, a factor N_s^2; the worked example gives $64\times$ at $N_s = 8$, $m = 100$.
- A parameter-count reduction for symmetry-constrained learning layers (derived in Section 4): circulant weight sharing reduces a cyclic layer from m^2 to m parameters, factor m; the worked example gives $65{,}536 \to 256$ parameters at $m = 256$.
- A thermodynamic framing of the gains (derived in Section 4): the Landauer bit energy $2.87098 \times 10^{-21}\,\mathrm{J}$ at $300\,\mathrm{K}$ and the corresponding $3.483 \times 10^{20}\,\mathrm{bits/J}$ ceiling, under the boundary conventions of [4].
- A falsification protocol and engineering pathway (Section 7): the smallest experiment that would refute the correspondence, stated as a proposal.

The sector-conservation result may also bear on problems outside physical computing wherever a learning system interacts with a structured physical process—edge inference on rotating machinery, and hybrid quantum-classical workflows whose classical driver must track symmetry sectors—as far as the derivations here support those analogies.

2. Background and Related Work

Symmetry reduction in hybrid dynamics. [3] studies simple hybrid systems $\mathbf{H} = (D, X, \mathcal{S}, \Delta)$, with a smooth vector field X on a manifold D, a co-dimension-one switching surface $\mathcal{S}$, and a reset map Δ, applied to bipedal walking and embedded computer systems. Its central technical asset for us is the hybrid Routhian construction: when the configuration space is $Q = P \times G$ with G an abelian Lie group, and both the Lagrangian L and a unilateral constraint function h are cyclic (invariant under the group action), the system reduces to a Routhian $R_c^\mu : TP \to \mathbb{R}$ and a reduced constraint $\bar{h} : P \to \mathbb{R}$. This is the exact mechanical precedent for our computational claim: invariance under a continuous group action removes degrees of freedom without loss of dynamical fidelity. Our Section 3 transposes the same group-action structure from TQ to a computational state space.

Physical substrates as computers. [9] formalizes the Inverted Church-Turing-Deutsch thesis, defining computation "not as the manipulation of symbols, but as the time evolution of physical systems governed by engineered Hamiltonians," and distinguishes Class A (Ising-based optimization, Hamiltonian H_{Ising}) from Class B (spectral synthesis, H_{Spectral}). Its methodology encodes problem sizes $N = 3$ to 12 spin variables with 20 independent replicates per size and treats the time-evolution operator $U(t) = e^{-iHt/\hbar}$ as the computational step. Our framework is complementary: it tells the Hamiltonian engineer which symmetry to build into H so that the resulting evolution is block-diagonal by construction. [9]'s Class B pipeline—inverse spectral engineering, translating a target eigenvalue list into a potential $V(x)$—is precisely where a cyclic symmetry constraint would cut the search dimensionality.

Hybrid orchestration and locality. [2] presents Tierkreis, a higher-order dataflow graph runtime for hybrid quantum-classical algorithms, motivated by the back-and-forth between classical and quantum phases that "makes understanding, debugging, and verifying quantum software unnecessarily difficult"; its nodes are pure functional primitives connected by edges, and its contribution is a semantics suited to distributed, asynchronous hybrid execution. [5] introduces the Quantum Execution Locality Framework (QELF), representing hybrid workflows as directed dataflows whose key characteristic is quantum execution locality—the extent to which computation stays resident on the QPU—and concluding that "the effectiveness of quantum acceleration depends not only on algorithmic capability, but also on the extent to which computation remains resident on the QPU." Both works frame the orchestration layer within which a symmetry-exploiting substrate would sit: a conserved sector label is exactly the kind of invariant metadata a dataflow runtime can track to avoid redundant synchronization, since states in one sector never need to communicate with states in another.

Thermodynamic accounting. [4] proposes two bits-per-joule metrics—Thermodynamic Epiplexity per Joule (new bits of structure about an environment variable encoded per unit energy) and empowerment per joule—under explicit accounting conventions separating measured energy E_{cons} inside a boundary from thermodynamic dissipation Q_{diss}. Its Proposition 1 is a decoupling construction showing why boundary closure is essential, and its Corollary 1 gives a closed-cycle Landauer-scale benchmark. We adopt its conventions in Section 4: our energy figures are stated inside an explicit boundary, and our efficiency claims are ratios against the Landauer bit energy rather than absolute intelligence scores, following [4]'s position that "comparisons are only meaningful under explicit boundary, coarse-graining/noise, horizon/reset, and cost conventions."

Edge deployment pressure. [1] surveys the architectures and enabling techniques for training and inference at the network edge, tabulating systems such as FedAvg, BlockFL and GossipGraD that address the privacy-driven constraint of distributed data, and cataloging inference systems whose configuration parameters must be optimized in real time. This is the demand side of our problem: symmetry-exploiting substrates would change the configuration space itself, replacing a high-dimensional tuning problem with a structured one.

Redundancy elimination as a computational primitive. [8] builds a formal system whose single primitive is the act of drawing a boundary, with a syntactic void as identity element, and states a reduction law of idempotence: "two adjacent and identical boundaries are structurally equivalent to a single boundary." Read computationally, idempotence is the simplest redundancy-elimination rule; symmetry is its continuous generalization, and our block decomposition is the general mechanism by which such redundancy becomes a resource rather than a nuisance. [11] argues a stronger and more speculative thesis—that circular topology is ontologically primary, with the circle manifold $S^1 = \{z \in \mathbb{C} \mid |z| = 1\} \cong \mathbb{R}/\mathbb{Z}$ as the minimal structure generating both continuity and discreteness—and supplies verification strategies designed to keep the framework falsifiable. We do not adopt [11]'s ontological claim; we use only its mathematical object, S^1, which is also our symmetry group's underlying manifold, and its insistence on experimental signatures as a model for our own falsification protocol in Section 7.

Evaluation and hypothesis selection. [6] evaluates approaches to artificial general intelligence and concludes a hybrid approach is necessary; its treatment of experimental design is directly relevant to our Section 7: "Each experiment costs resources, and the information gained should be evaluated in terms of expected benefit across hypotheses and future experiments," with a recommendation toward geometric-mean utility when future experiments depend on preceding outcomes. Our falsification protocol is designed as the cheapest experiment that discriminates between the symmetry-constrained and unconstrained hypotheses. [7] surveys tests of machine intelligence, noting that Turing's imitation game "has attracted much criticism" because a machine could appear intelligent without real intelligence, and cataloging compression tests in the style of Mahoney, in which comprehension is estimated by predicting missing words from a passage. [7]'s summary criteria—valid, informative, wide range, general, dynamic, unbiased, fundamental—supply the rubric we apply in Section 6 to the question of what our efficiency metric actually measures. [10] establishes a claimed categorical equivalence between biological and engineered computational systems, including a reflexive-domain construction $D_\infty = \varprojlim D_n$ with $D_{n+1} = [D_n \to D_n]$ and a fixed-point combinator, and proposes a multi-tiered validation framework spanning two-dimensional electronic spectroscopy, neuromorphic benchmarking, and hybrid bio-artificial systems. Its validation tiering is a template for how our own claims should be tiered, and its energy-management isomorphisms are the biological analogue of the conservation structure we derive for engineered substrates.

3. Methods

3.1 Model of a physical computing substrate

We model a substrate as a triple $\Sigma = (\mathcal{H}, U, \mathcal{M})$:

- $\mathcal{H}$ is the state space. For a classical linear substrate (a network of resonators, a cyclic truss, a coupled-oscillator ring) $\mathcal{H} = \mathbb{R}^M$ with the Euclidean inner product; for a coherent substrate $\mathcal{H}$ is an M-dimensional Hilbert space, matching the setting of [9] where the evolution operator is $U(t) = e^{-iHt/\hbar}$.
- $U : \mathcal{H} \to \mathcal{H}$ is the update map: for classical substrates a symmetric matrix operator A applied per step or per mode solve; for coherent substrates the Hamiltonian propagator.
- $\mathcal{M} : \mathcal{H} \to \mathcal{Y}$ is the readout map that extracts a computational answer.

This abstraction is deliberately hardware-neutral, in the same spirit as the dataflow abstraction of [5], which represents hybrid applications as directed dataflows "intentionally independent of programming models or hardware implementations."

3.2 Symmetry action and the computational Noether correspondence

Let $G = SO(2)$ act unitarily on $\mathcal{H}$ through a representation $\rho : G \to \mathrm{U}(\mathcal{H})$. Over $\mathbb{C}$, every unitary representation of $SO(2) \cong U(1)$ decomposes into one-dimensional irreducible representations (characters) $\chi_k(\theta) = e^{ik\theta}$, $k \in \mathbb{Z}$; over $\mathbb{R}$ the irreducibles are the two-dimensional rotation blocks with harmonic index $k \geq 1$ plus the invariant subspace $k = 0$. Define the symmetry-adapted decomposition

$$\mathcal{H} = \bigoplus_k \mathcal{H}_k, \qquad \mathcal{H}_k = \{v \in \mathcal{H} \;\mid\; \rho(\theta)v = e^{ik\theta}v \text{ for all } \theta\}.$$

Proposition 1 (Noether-type correspondence for computation). If the update map commutes with the group action, $U \circ \rho(\theta) = \rho(\theta) \circ U$ for all $\theta \in SO(2)$, then:

1. (Conservation) Each subspace $\mathcal{H}_k$ is invariant under U: a state's harmonic index k—its computational charge—is conserved by the dynamics.
2. (Block decomposition) $U = \bigoplus_k U_k$ with $U_k = U|_{\mathcal{H}_k}$, so any function computed by iterating or diagonalizing U decomposes into independent per-sector computations.

3. (No readout mixing) Any readout $\mathcal{M}$ that is itself symmetry-equivariant maps sector k inputs to sector-k-labeled outputs; sector labels can be read off without cross-sector computation.

Proof. (1) For $v \in \mathcal{H}_k$ and any θ, $\rho(\theta)Uv = U\rho(\theta)v = e^{ik\theta}Uv$, so $Uv \in \mathcal{H}_k$. (2) Since each $\mathcal{H}_k$ is invariant and the decomposition is orthogonal, U is block diagonal in the symmetry-adapted basis. (3) Equivariance of $\mathcal{M}$ gives $\mathcal{M}(\mathcal{H}_k) \subseteq \mathcal{Y}_k$ where $\mathcal{Y}_k$ is the sector-k output subspace. □

This is the computational analogue of the mechanical statement in [3]: there, a cyclic Lagrangian yields a conserved momentum map and a reduced Routhian $R_c^\mu : TP \to \mathbb{R}$ on the shape space P; here, a commuting update map yields a conserved sector label and reduced per-sector operators U_k. The hybrid structure of [3]—with a guard $\mathcal{S}$ and reset map Δ—also transposes: if the reset map is symmetry-equivariant ($\Delta \circ \rho(\theta) = \rho(\theta) \circ \Delta$ on $\mathcal{S}$), the conservation law survives the discrete events, so the correspondence extends to hybrid (continuous-plus-discrete) computation exactly as [3] extends Routh reduction to hybrid Lagrangian systems.

3.3 Cost model

We use a standard dense linear-algebra cost model. A dense eigensolve or matrix multiply on an $n \times n$ operator costs $c\,n^3$ floating-point operations with a technology-dependent constant c; a fast transform (DFT) between the physical and symmetry-adapted bases costs $O(n \log n)$. All speedup ratios below are independent of c. For learning layers we count parameters, not operations. Energy accounting follows [4]: we report ratios against the Landauer bit energy inside an explicit boundary, and we do not claim that any physical substrate approaches that bound.

3.4 What we compute and what we assume

Section 4 derives four quantities by explicit arithmetic: the sector speedup N_s^2, a worked instance of it; the parameter-sharing reduction m with a worked instance; the Landauer bit energy and bits-per-joule ceiling at $T = 300\,\mathrm{K}$; and an illustrative harmonic-truncation budget under a stated spectral-concentration assumption. No experiment, simulation, or measurement is performed in this paper; the validation experiment is specified as a proposal in Section 7, following the resource-aware experimental-design stance of [6].

4. Analysis

4.1 Sector decomposition speedup

Consider a cyclic substrate with N_s identical sectors, total dimension $M = N_s\,m$. Without exploiting symmetry, a dense spectral solve costs

$$C_{\text{full}} = c\,M^3 = c\,(N_s m)^3 = c\,N_s^3\,m^3.$$

With Proposition 1, the solve decomposes into N_s independent solves of size m (one per harmonic sector), plus a symmetry-adapted change of basis costing $O(M \log M)$, which is lower order:

$$C_{\text{sym}} = N_s\,c\,m^3 + O(M \log M).$$

The speedup ratio is therefore

$$S = \frac{C_{\text{full}}}{C_{\text{sym}}} \approx \frac{c\,N_s^3 m^3}{N_s\,c\,m^3} = N_s^2.$$

Worked example. Take $N_s = 8$ sectors of size $m = 100$, so $M = 800$. Then

$$C_{\text{full}} = c \cdot 800^3 = c \cdot 5.12 \times 10^8,$$

$$C_{\text{sym}} = 8 \cdot c \cdot 100^3 = 8 \cdot c \cdot 10^6 = c \cdot 8.0 \times 10^6,$$

$$S = \frac{5.12 \times 10^8}{8.0 \times 10^6} = 64 = N_s^2.$$

The transform overhead is $O(800 \log 800) \approx O(800 \times 9.64) \approx 7.7 \times 10^3$ units, negligible against 8.0×10^6 (a fraction $\approx 9.6 \times 10^{-4}$), so the $64\times$ figure stands. This is the cyclic-structure analogue of reducing a structural analysis to a single sector, and it is exact, not asymptotic: the block decomposition of Proposition 1 is an identity.

4.2 Parameter reduction in symmetry-constrained learning layers

A generic linear layer acting within one sector on an m-dimensional sector space carries a dense matrix $W \in \mathbb{R}^{m \times m}$: m^2 parameters. If the layer is constrained to be circulant—equivariant under the cyclic group $\mathbb{Z}_m$, the finite analogue of the $SO(2)$ action—then W is diagonalized by the DFT matrix F_m,

$$W = F_m^{-1} \, \mathrm{diag}(\hat{w}_0, \ldots, \hat{w}_{m-1}) \, F_m,$$

and is fully specified by the m eigenvalues $\hat{w}_k$. The parameter count falls from m^2 to m, a reduction factor

$$R = \frac{m^2}{m} = m.$$

Worked example. At $m = 256$:

$$m^2 = 256^2 = 65{,}536 \quad \longrightarrow \quad m = 256,$$

$$R = \frac{65{,}536}{256} = 256.$$

The same sharing multiplies through depth: a network of L such layers drops from $L\,m^2$ to $L\,m$ parameters, so at $L = 12$, $m = 256$: from $12 \times 65{,}536 = 786{,}432$ to $12 \times 256 = 3{,}072$, again a factor 256. This is the symmetry-constrained-autoencoder hypothesis made precise: the constraint is not a regularizer but an exact equivariance condition, and its cost is the restriction of expressiveness to permutation-circulant maps.

4.3 Landauer floor and bits-per-joule ceiling

The Landauer bound for erasing one bit at temperature T is

$$E_{\text{bit}} = k_B \, T \, \ln 2,$$

with $k_B = 1.380649 \times 10^{-23}\,\mathrm{J/K}$. At $T = 300\,\mathrm{K}$:

$$k_B \, T = 1.380649 \times 10^{-23} \times 300 = 4.141947 \times 10^{-21}\ \mathrm{J},$$

$$E_{\text{bit}} = 4.141947 \times 10^{-21} \times 0.693147 = 2.87098 \times 10^{-21}\ \mathrm{J}.$$

The corresponding ceiling on learning efficiency, in the sense of the bits-per-joule metrics of [4], is

$$\eta_{\max} = \frac{1}{E_{\text{bit}}} = \frac{1}{2.87098 \times 10^{-21}} = 3.483 \times 10^{20}\ \mathrm{bits/J}.$$

For scale, a system dissipating $P = 20\,\mathrm{W}$—the figure [4] gives for the human brain—has a Landauer-limited throughput ceiling of

$$\frac{P}{E_{\text{bit}}} = \frac{20}{2.87098 \times 10^{-21}} = 6.966 \times 10^{21}\ \mathrm{bits/s}.$$

These are ceilings, not achievements: [4] is explicit that measured consumption E_{cons} and dissipation Q_{diss} must be distinguished and that boundary closure is essential, and we adopt that discipline. The role of these numbers in our framework is to convert the operation-count savings of Section 4.1 and the parameter savings of Section 4.2 into a common currency: every operation and every parameter not needed is energy not spent, and the headroom between any realized substrate and $\eta_{\max}$ measures how much of the thermodynamic budget is being spent on symmetry-ignorant computation.

4.4 Harmonic truncation budget

If a cyclic substrate's task-relevant signal concentrates its energy in the lowest K of N_s harmonic sectors, then computing only those sectors costs a fraction K/N_s of the full symmetric cost. We state the required concentration as an explicit assumption, since no measurement of a task spectrum is available in this paper: **we assume** a task distribution for which the first $K = N_s/4$ harmonics carry at least 90% of the signal energy. Under that assumption the truncation budget is

$$\frac{K}{N_s} = \frac{N_s/4}{N_s} = 0.25,$$

so the compute cost is $4\times$ lower than the full symmetric solve (which is itself N_s^2 lower than the dense solve), while at most 10% of signal energy is discarded. Combining with Section 4.1 at $N_s = 8$:

$$S_{\text{total}} = N_s^2 \times \frac{N_s}{K} = 64 \times 4 = 256\times$$

relative to the dense solve, under the stated assumption. The assumption is the load-bearing element: if the task energy is spread evenly across sectors ($K = N_s$), the truncation factor is 1 and only the exact N_s^2 saving survives. Section 6 treats this sensitivity as a primary limitation.

4.5 Conservation as orchestration metadata

Proposition 1(1) has an orchestration consequence that connects to the dataflow literature. In a hybrid workflow of the kind [2] and [5] describe, the classical driver must decide when results from the physical device can be combined, post-processed, or reused. If the device conserves sector labels, then two intermediate states with different labels can never interact, so the driver may process sectors independently and in parallel without a consistency check; the conserved quantity plays the role that type information plays in a dataflow graph. In [2]'s terms, the sector label is a pure functional invariant of the primitives; in [5]'s terms, it enlarges the contiguous compute region—the extent to which computation stays resident on the accelerator—because cross-sector synchronization is provably unnecessary. We quantify this only structurally (the synchronization graph factorizes into N_s disjoint components) and leave the end-to-end latency accounting to future work, since no source gives measured synchronization overheads for a specific hybrid stack.

5. Results

All numbers below are computed in Section 4 by the arithmetic shown there; none is measured or simulated.

1. **Exact sector speedup.** A substrate with N_s equivalent sectors under an $SO(2)$-commuting update map solves the spectral problem N_s^2 times faster than the dense method (Section 4.1). Worked instance: $N_s = 8$, $m = 100$, $M = 800$: $C_{\text{full}} = c \cdot 5.12 \times 10^8$ versus $C_{\text{sym}} = c \cdot 8.0 \times 10^6$, speedup $64\times$; the basis-transform overhead is $\approx 9.6 \times 10^{-4}$ of the reduced cost and does not change the figure.

1. **Parameter reduction.** Circulant (cyclic-equivariant) weight sharing reduces a layer from m^2 to m parameters, factor m (Section 4.2). Worked instance: $m = 256$: $65{,}536 \to 256$ parameters per layer; over $L = 12$ layers, $786{,}432 \to 3{,}072$, factor 256.

1. **Landauer floor.** $E_{\text{bit}} = k_B T \ln 2 = 2.87098 \times 10^{-21}\,\text{J}$ at $T = 300\,\text{K}$ (Section 4.3), giving a ceiling $\eta_{\max} = 3.483 \times 10^{20}\,\text{bits/J}$, and a 20 W dissipation ceiling of $6.966 \times 10^{21}\,\text{bits/s}$.

1. **Truncation budget (projection under a stated assumption).** If the lowest $K = N_s/4$ harmonics carry $\geq 90\%$ of task signal energy (assumption; no measured spectrum is available in this paper), harmonic truncation cuts the symmetric cost by a further factor 4, for a combined $256\times$ over the dense solve at $N_s = 8$ (Section 4.4). Uncertainty: the factor degrades linearly to $64\times$ as $K \to N_s$; the combined figure holds only within the stated energy-concentration assumption.

1. **Structural result (non-numeric).** The conserved sector label of Proposition 1 factorizes the synchronization graph of a hybrid workflow into N_s disjoint components (Section 4.5), a qualitative but exact consequence used in Section 7.

6. Discussion

What the theorem does and does not say. Proposition 1 is exact linear algebra: commuting update maps are block diagonal, full stop. The quantitative gains, however, depend on three contingent conditions: that the substrate's useful dynamics genuinely commute with a large continuous symmetry; that the task's information content is concentrated in few sectors; and that the readout is equivariant. A substrate with weak or broken symmetry gains little; a task whose signal spreads evenly over sectors gains only the exact N_s^2 factor and nothing from truncation; a readout that mixes sectors forces cross-sector computation at the output, eroding the savings. The framework's value is therefore as a design compass—maximize exploitable symmetry, then check sector concentration—rather than as a universal speedup guarantee.

Relation to the mechanical precedent. [3]'s hybrid Routhian reduction requires both the Lagrangian and the constraint function to be cyclic; our Proposition 1(1) similarly requires equivariance of both the continuous update and the discrete reset. The interesting open question is what happens when the reset map Δ is only approximately equivariant —e.g., a manufacturing defect that breaks one sector's equivalence. In the mechanical setting [3] treats, symmetry breaking destroys the reduction; in the computational setting, a small equivariance defect should leak a bounded amount of amplitude between sectors at a rate set by the defect size, but we have not derived that leakage bound here, and it is the most important open analysis.

Limitations list. The results of Sections 4 and 5 rest on the following assumptions, each with its sensitivity stated:

- **Equivariance of the update map.** Proposition 1 requires $U \circ \rho(\theta) = \rho(\theta) \circ U$ exactly. The N_s^2 speedup is exact under this condition and vanishes entirely if the symmetry is absent; there is no graceful degradation derived here. Sensitivity: total.
- **Equivariance of the reset map.** As argued above, the extension to hybrid computation assumes $\Delta \circ \rho(\theta) = \rho(\theta) \circ \Delta$ on $\mathcal{S}$. A defect in one sector is not analyzed; the leakage bound is open. Sensitivity: unknown, plausibly linear in defect size, but undemonstrated.
- **Sector concentration of the task.** The $256\times$ combined figure of Section 4.4 holds only under the assumed $\geq 90\%$ energy in the lowest $K = N_s/4$ harmonics. Sensitivity: linear; the truncation factor degrades from 4 to 1 as $K \to N_s$, leaving the exact $64\times$.
- **Equivariant readout.** If $\mathcal{M}$ mixes sectors, the output stage reintroduces cross-sector computation. Sensitivity: proportional to the readout's off-diagonal weight.
- **Dense cost model.** The $c\,n^3$ model ignores structure-aware dense solvers (e.g., banded or low-rank methods) that could reduce C_{full} without symmetry; the speedup against such a solver would be smaller. Sensitivity: the ratio S is an upper bound relative to the best structure-blind baseline.

- **Landauer framing.** The bits-per-joule figures are ceilings under the boundary conventions of [4]; no claim is made that any substrate approaches them, and the operation-count savings of Section 4.1 are not converted into joules for any specific device.

What would falsify the main result. The correspondence itself (Proposition 1) is a theorem of linear algebra and cannot be falsified except by an error in its proof; what is falsifiable is its *usefulness claim*—that exploitable continuous symmetry is a decisive design variable for physical computing substrates. The falsifying observation is concrete: build the symmetry-constrained autoencoder of Section 7 and find that its reconstruction fidelity at equal parameter count is *worse* than an unconstrained baseline of the same parameter budget, or that a cyclic substrate with a deliberately broken sector shows no measurable cost or fidelity difference from the fully symmetric one. Either outcome would show that the conserved sector label carries no computational value in practice, and the framework would reduce to a restatement of block diagonalization with no design content.

Relation to evaluation criteria. [7]'s summary rubric—valid, informative, wide range, general, dynamic, unbiased, fundamental—applies uncomfortably to our efficiency metric. Bits-per-joule headroom against $\eta_{\max}$ is informative (a scalar) and fundamental (tied to a physical constant), but it is not valid as a measure of intelligence: it measures resource discipline, not capability, exactly the gap [7] identifies when it notes that a test "should capture intelligence and not some related quantity." We therefore present the metric as a *substrate-selection* metric, not an intelligence metric. [6]'s experimental-design stance—that "each experiment costs resources, and the information gained should be evaluated in terms of expected benefit across hypotheses and future experiments"—is satisfied by choosing the autoencoder comparison as the first test, since it simultaneously probes the parameter-reduction claim (Section 4.2) and the fidelity-preservation claim at minimal cost. [10]'s multi-tiered validation framework, spanning spectroscopic measurement, neuromorphic benchmarking, and hybrid systems, supplies the longer-term tiering: our Section 7 proposal sits at its lowest tier, and the substrate-level claims would graduate upward only through hardware demonstration.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are designers of physical computing substrates—engineers specifying resonator networks, cyclic mechanical structures, or engineered Hamiltonians in the sense of [9]—and architects of hybrid orchestration layers of the kind [2] and [5] build. For the substrate designer, the result changes the geometry decision: when a choice exists between a generic coupled network and a cyclically symmetric one of equal task performance, the symmetric choice carries a provable N_s^2 spectral-solve advantage and a factor-m parameter advantage in any learned layer wrapped around it, so symmetry becomes a quantified selection criterion rather than an aesthetic preference. For the orchestration architect, the conserved sector label becomes trackable invariant metadata: a runtime in the style of [2], whose nodes are pure functional primitives connected by edges, can schedule per-sector subgraphs independently, and in [5]'s locality vocabulary the label enlarges the contiguous compute region by certifying that cross-sector synchronization is unnecessary. For edge-system designers facing the high-dimensional configuration problem [1] documents—frame rate, resolution, model selection, early-exit tuned in real time—a symmetry-structured model family shrinks the configuration space itself, because sector membership is fixed by the substrate rather than learned.

(b) Readiness level. This work is analytical and unreviewed: the correspondence is derived, the cost model is computed on paper, and no hardware or software artifact exists. Against the NASA Technology Readiness Level definitions, it sits at TRL 1–2 (basic principles observed and reported; concept formulated). It cannot exceed TRL 3 without an analytical or experimental proof of concept on a concrete substrate, and no manufacturing readiness is claimed.

(c) The smallest next test. The minimal falsifying experiment is a software-only study, no physical device required: train two autoencoders on a task with known cyclic structure (e.g., rotation-equivariant image or sensor data), one with circulant-constrained layers of the Section 4.2 form (m parameters per layer, $W = F_m^{-1}\,\mathrm{diag}(\hat{w})\,F_m$) and one uncon-

strained baseline matched to the *same total parameter budget* by width adjustment. The study would measure reconstruction error (e.g., mean squared error) and training energy per epoch for both. The hypothesis is refuted if the constrained model's fidelity is worse at equal parameter budget, or if its energy-per-epoch advantage fails to track the predicted factor-m parameter reduction within a stated tolerance. A second, cheap calculation would de-risk the truncation claim of Section 4.4: compute the harmonic energy spectrum E_k of a candidate task's data across sectors and check whether the lowest $K = N_s/4$ harmonics carry $\geq 90\%$ of the energy; a flat spectrum refutes the $256\times$ projection for that task while leaving the exact N_s^2 result intact.

(d) Binding constraints. In order of expected severity, as far as this paper's derivations support them: first, symmetry fidelity of the physical device—manufacturing tolerances that break sector equivalence erode the conservation law, and the leakage rate is the open analysis of Section 6; second, readout design, since an equivariant readout is required for the savings to survive the output stage; third, task-spectrum concentration, which is a property of the data, not the device, and must be verified per application; fourth, the expressiveness cost of circulant constraint, which restricts learned layers to permutation-circulant maps and may dominate the parameter savings on tasks without cyclic structure. Whether fabrication tolerances can hold sector equivalence at the defect sizes the leakage bound (once derived) permits is an open question this paper cannot settle.

(e) Order-of-magnitude bounds (projections). Under the assumptions already stated—dense cost model, $N_s = 8$, $m = 100$, and the 90% low-harmonic concentration—the combined compute reduction is $256\times$ over a dense solve (Section 4.4). As a labeled projection with its assumptions: if a substrate's spectral-solve energy were proportional to operation count at fixed *c*, the same factor would apply to that energy component; no absolute joule figure is derivable here because no source gives a measured per-operation energy for any specific substrate, and we do not invent one. The only absolute energy number in this paper is the Landauer ceiling $E_{\text{bit}} = 2.87098 \times 10^{-21}\,\text{J}$ and its derived throughput ceilings, which are bounds, not performance projections.

(f) Risks and misuse paths. The main misuse is rhetorical: quoting the $256\times$ or $64\times$ figures as device-level energy savings without the equivariance, concentration, and readout conditions, or quoting the Landauer ceiling as if it were an achievable efficiency. Both figures are conditional on assumptions stated in Sections 4.4 and 6, and the ceiling is a thermodynamic bound of the kind [4] explicitly frames as convention-dependent. A secondary risk is over-constraining learned models: applying circulant constraints to tasks without cyclic structure would degrade fidelity while delivering no conservation benefit, since the equivariance condition would be violated by the data rather than the device.

8. Conclusion

We have transported the symmetry-reduction structure of hybrid mechanics into physical computing. The central result, Proposition 1, states that any substrate whose update map commutes with an $SO(2)$ action conserves a computational quantity—the harmonic sector label—and decomposes into independent per-sector blocks, extending to hybrid computation when the reset map is equivariant, in direct analogy with the hybrid Routhian reduction of [3]. From this conservation law we derived, with explicit arithmetic: an exact N_s^2 spectral-solve speedup ($64\times$ at $N_s = 8$, $m = 100$); a factor-m parameter reduction for circulant learning layers ($65{,}536 \to 256$ at $m = 256$, and $786{,}432 \to 3{,}072$ over twelve layers); a conditional $256\times$ truncation budget under a stated 90% low-harmonic energy assumption; and the Landauer-scale energy framing $E_{\text{bit}} = 2.87098 \times 10^{-21}\,\text{J}$ at $300\,\text{K}$, ceiling $\eta_{\max} = 3.483 \times 10^{20}\,\text{bits/J}$, under the boundary conventions of [4]. The design rule that follows—maximize exploitable continuous symmetry, then verify sector concentration and readout equivariance—converts substrate geometry from a matter of craft into a quantified choice, and gives dataflow runtimes of the kind [2] and [5] describe a provably synchronizable decomposition. The framework's honest scope is narrow: it guarantees nothing for symmetry-broken devices, sector-spread tasks, or mixing readouts, and the leakage bound under equivariance defects is the most important open analysis. The smallest falsifying test—circulant-constrained versus parameter-matched unconstrained autoencoders—is specified in Section 7 and is ex-

ecutable without new hardware. Beyond physical computing, the sector-conservation structure may inform any setting where a learning system drives a symmetry-structured physical process, including edge inference on rotating machinery and classical control of symmetry-sectored quantum workflows, as far as the derivations here support those analogies.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.55.12). 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, openai/gpt-oss-120b, moonshotai/kimi-k2.6, 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: the execution sandbox was unavailable, so a second language model recomputed 11 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 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] Edge Intelligence: Paving the Last Mile of Artificial Intelligence with Edge Computing. arXiv:1905.10083v1. https://arxiv.org/abs/1905.10083v1 [2] Tierkreis: A Dataflow Framework for Hybrid Quantum-Classical Computing. arXiv:2211.02350v1. https://arxiv.org/abs/2211.02350v1 [3] Symmetries and periodic orbits in simple hybrid Routhian systems. arXiv:2001.08941v1. https://arxiv.org/abs/2001.08941v1 [4] Thermodynamic Limits of Physical Intelligence. arXiv:2602.05463v2. https://arxiv.org/abs/2602.05463v2 [5] Dataflows and Computational Patterns for Hybrid Quantum-Classical Scientific Computing. arXiv:2608.19348v1. https://arxiv.org/abs/2608.19348v1 [6] The Artificial Scientist: Logicist, Emergentist, and Universalist Approaches to Artificial General Intelligence. arXiv:2110.01831v1. https://arxiv.org/abs/2110.01831v1 [7] Tests of Machine Intelligence. arXiv:0712.3825v1. https://arxiv.org/abs/0712.3825v1 [8] https://papers.qnfo.org/papers/syntactic-generation-primitive-distinctions. QNFO: Syntactic Generation. [9] https://papers.qnfo.org/papers/programming-the-vacuum. QNFO: Programming the Vacuum. [10] https://papers.qnfo.org/papers/biological-and-engineered-computational-systems. QNFO: Biological and Engineered Computational Systems. [11] https://papers.qnfo.org/papers/winding-numbers-and-strange-loops. QNFO: WINDING NUMBERS AND STRANGE LOOPS.