V1.0.0

Noether-Bounded Computation: Conservation Laws as Native Constraints on Physical Processors

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Digital processors that emulate continuous dynamics pay a per-operation rounding tax: each step can violate a conserved quantity of the emulated system, and the violation accumulates with operation count. Whether a physical computing substrate whose operations directly instantiate a continuous symmetry group inherits the corresponding Noether conserved quantities as built-in constraints — and what advantage that inheritance buys in error and energy — has not been formally quantified. This paper derives a Noether-based cost model relating the symmetry group G of an analog processor to bounds on conservation-law drift and error growth. For a b-bit digital emulator we show the worst-case relative drift of a conserved quantity after N steps is $\Delta Q/Q \leq Nu$ with $u = 2^{-b}$, giving 2.33×10^{-4} at $N = 10^6$ steps in 32-bit mantissa arithmetic; for a G-equivariant continuous processor the drift is zero by Noether's theorem up to physical dissipation $\omega T/Q_{\text{res}}$. The model predicts a crossover operation count $N^* = (\omega T/Q_{\text{res}})/u \approx 4.30 \times 10^6$ steps under stated assumptions, beyond which the symmetry-aligned substrate has strictly lower conservation drift. Each independent conserved quantity also contracts the accessible error-state volume by a factor u. These bounds turn symmetry choice into a designable error-control parameter for analog and neuromorphic accelerator design.

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.). Noether-Bounded Computation: Conservation Laws as Native Constraints on Physical Processors. QNFO Preprint v1.0.0.

1. Introduction

Engineers who simulate oscillators, wave systems, and fluid-like dynamics on digital hardware face a structural problem, not a tuning problem. Every floating-point operation rounds, and each rounding event can push the computed state off the level set of a quantity that the modeled physical system conserves exactly — energy, momentum, circulation, or a norm. The accumulated drift is not random noise that averages away in all regimes: in worst case it grows linearly with the number of operations, and in long runs it can dominate the result. This is the same difficulty the earliest numerical analysts identified when fixed word lengths made machine arithmetic diverge from idealized arithmetic, with rounding errors in large matrix calculations able to accumulate and potentially dominate the answer [8]. It is also the difficulty behind the reproducibility concerns that now accompany computational modeling across the sciences [8]. The cost is concrete: longer runs require either more bits, compensated summation, structure-preserving integrators, or periodic projection — each an engineering overhead paid per computation.

There is a second, physical route. A processor whose continuous-time dynamics are exactly equivariant under a symmetry group G — a network of photonic resonators, an oscillator network, a symplectic analog substrate — does not approximate the conservation laws attached to that symmetry; it instantiates them. Noether's theorem connects symmetries of an invariant action to conserved currents, with the multipliers of the conservation laws given by the variational symmetries themselves [1]. If the hardware's Lagrangian has the symmetry, the conserved quantity is a constraint on every trajectory the hardware can execute, regardless of how many "operations" the trajectory implements. The open

question, posed in the research program this paper belongs to, is whether one can formally quantify an efficiency or error-correction advantage for such symmetry-equivariant physical processors — for example, that respecting a conserved quantity bounds dissipation or error accumulation per operation.

This paper answers a restricted but sharp version of that question. We build a cost model that compares two implementations of the same continuous computation: (i) a digital emulator that discretizes the dynamics and rounds at each step, and (ii) a G-equivariant continuous processor whose dynamics conserve the Noether quantities of the emulated system identically. The model yields three quantitative results. First, a drift bound: the emulator's conserved-quantity error after N steps is bounded by Nu (worst case) or $u\sqrt{N}$ (unbiased-rounding random walk), where $u = 2^{-b}$ is the unit roundoff, while the equivariant processor's drift is zero up to a physical dissipation term set by its quality factor. Second, a crossover: there is an operation count N^* beyond which the symmetry-aligned substrate has strictly lower conservation drift, and we compute N^* in closed form and numerically under stated assumptions. Third, a state-space contraction: each independent conserved quantity removes one error direction from the processor's phase space, contracting the volume of reachable error states by a factor u per conserved quantity.

If the model holds, the decision it changes is architectural. Designers of analog accelerators, photonic oscillator networks, and neuromorphic systems currently choose network structure for function, not for the conservation algebra the structure enforces; the model says the symmetry group of the substrate is an error-control design parameter, on the same footing as word length in digital design. Numerical analysts gain a complementary statement: the drift bounds quantify exactly what structure-preserving integrators buy and what remains unpaid. The results may also bear on adjacent problems where conserved quantities discipline long computations, such as long-horizon molecular or geophysical simulation, insofar as those computations conserve quantities of the types classified for volume, surface, and curve domains in three dimensions [2].

The contributions are:

- A Noether drift theorem for emulated versus instantiated conservation laws: worst-case drift Nu and random-walk drift $u\sqrt{N}$ for a b-bit emulator versus zero-plus-dissipation for a G-equivariant processor, derived in Sections 4.1–4.2.
- A closed-form crossover operation count $N^* = (\omega T/Q_{\text{res}})/(Cu)$ with a worked numerical value $N^* \approx 4.30 \times 10^6$ under stated assumptions, derived in Section 4.3.
- An error-state volume contraction of factor u per independent conserved quantity, with the $r = 1$ and $r = 2$ cases computed, derived in Section 4.4.
- A mapping from the conservation-law literature to processor design constraints — which invariants can be enforced locally, which arise only for special equation classes, and what partial symmetry buys — given in Sections 2 and 5.

The analysis speaks to adjacent fields in which conserved quantities discipline long computations — long-horizon geophysical and molecular simulation, and analog signal processing — insofar as those computations rely on invariants of the types classified in [2].

2. Background and Related Work

The mathematical side of this paper rests on the theory of conservation laws for differential equations, and the computational side on the known failure modes of finite-precision emulation. We review both, citing each work for the specific result its text establishes.

The multiplier framework for conservation laws is the formal backbone of our drift theorem. For a field equation, local conservation laws arise through multipliers, with the product of multiplier and equation a total divergence, and when the equation possesses a Lagrangian, Noether's theorem shows the multipliers are exactly the variational symmetries of

the action [1]. That work also gives an algebraic conservation law formula for scaling-invariant equations that generates all local conservation laws with non-zero scaling weight, applied to soliton equations such as Korteweg–de Vries, whose scaling $t \to \lambda^3 t$, $x \to \lambda x$, $u \to \lambda^{-2} u$ it records explicitly [1]. We use two things from it: the Noether identification of symmetries with multipliers (which licenses "symmetry of the hardware Lagrangian implies conserved quantity of the hardware trajectory"), and the fact that conservation laws can be generated algebraically from group structure (which makes the symmetry-to-invariant map computable at design time, not discovered per system).

Which conservation laws a physical processor can enforce depends on the dimensionality of the domain on which they hold. A unified study of PDE systems in three spatial dimensions distinguishes dynamical conservation laws on volumes, surfaces, and curves from topological conservation laws on surfaces and curves, gives local continuity-equation forms in which a density's time derivative is balanced by the divergence of a flux, and characterizes when two conservation laws are equivalent and when a law is trivial [2]. That classification matters for hardware: a volume-type law (energy, charge) can be enforced by local balance in a lumped-element network, while curve-type circulation laws require the substrate to realize the corresponding flux topology. The same study's distinction between dynamical, constraint, and constitutive equations [2] maps onto our separation of computational dynamics, conservation constraints, and physical dissipation.

Conservation is a structural property, not a generic one, and the multi-peakon literature makes this precise in a way directly relevant to design. For the general class $m_t + f(u, u_x)m + (g(u, u_x)m)_x = 0$, a classification determines exactly which families of these nonlinear dispersive equations possess conserved momentum, conserved H^1 norm, conserved H^2 norm, or conserved L^2 norm of m, with the conserved families given in closed form — for example the family $m_t + 2u_x h(u, u_x)m + u(h(u, u_x)m)_x = 0$ conserving the H^1 norm for arbitrary smooth h [3]. The lesson we import: two substrates that differ by one structural coefficient can differ by an entire conserved quantity, so the symmetry/structure audit of a processor design is not a formality but determines which error directions exist at all.

Integrable systems show the strongest version of the phenomenon. For the nonlocal two-dimensional periodic Toda lattice, the Lax-pair representation with quantized spectral parameter generates, via Darboux transformations and expansions in powers of the spectral parameter, hierarchies of infinitesimal symmetries and infinitely many conservation laws, and the multi-soliton solutions depend on infinitely many free parameters [4]. An infinitely many-conservation-law system is the limit case of our model: a substrate realizing such dynamics has an infinite tower of exact constraints, leaving correspondingly few error directions. We use this as the upper anchor of the "conservation richness" scale our cost model sits on.

Symmetry reduction is also a computational act. For the geopotential forecast equation, group foliation theory moves from the original equation to automorphic and resolving systems written in differential invariants, and invariant solutions obtained this way can be more general than those from the standard symmetry method; for the case $\beta \neq 0$ the resolving system is consistent and in involution [5]. We read this as a precedent for our central design move: computing in symmetry-adapted variables replaces a constraint-violation problem with a smaller exact problem. The geopotential equation itself — used for forecasting geopotential of mid-level vortices in the atmosphere [5] — is also a reminder that the continuous problems at stake are operational, not toy.

Conservation laws constrain even stochastic and irreversible dynamics. A stochastic collapse equation built on correlating interactions is consistent with exact conservation of momentum and orbital angular momentum in individual experiments, and with energy conservation where conventional nonrelativistic quantum theory predicts it, without introducing new physical constants, because collapse strength and duration are set by the ratio of interaction energy to total relativistic energy [6]. This is the closest physical precedent to our claim: a physical process can be fundamentally nonlinear and stochastic yet conservation-exact, so conservation-exactness is achievable in real interacting substrates rather than being a property of linear idealizations only.

Partial symmetry also pays. The theory of sub-symmetries — symmetries of a sub-system on solutions of the whole system — proves a theorem associating to each sub-symmetry of a quasi-Noether system a local conservation law, and for the Euler equations generates infinite conservation-law families including the Casimir invariants of two-dimensional flow, generalized momenta and angular momenta, and generalized helicity [7]. For processor design this means an incompletely symmetric substrate is not unprotected: each sub-symmetry still purchases a conserved quantity and thus removes one error direction, a graded version of our u^r contraction.

On the computational side, the historical record establishes the baseline we measure against. The pioneers of electronic computation recognized immediately that fixed word lengths make machine arithmetic differ from idealized arithmetic; von Neumann and Goldstine's 1947 analysis of matrix inversion explicitly treated the potential for rounding errors to accumulate and dominate large-scale calculations [8]. The same historical account records the alternative-arithmetics response — interval arithmetic guaranteeing that each computed quantity is a bracketing interval, and other paradigms that have not displaced IEEE 754 because of their own trade-offs in performance, applicability, or complexity [8] — which frames our proposal as a different axis of escape: not better number representations, but dynamics that make the quantity exact. The chaos-theory record, in which Lorenz's rounded restart of a three-ODE convection model produced divergent long-term behavior purely from finite precision [8], supplies the qualitative motivation for bounding drift at the substrate level. The reproducibility pressures of modern computational science [8] supply the demand side.

Two works sit closer to this paper's program. A categorical-equivalence framework for biological and engineered computational systems proposes that molecular and engineered devices implement categorically equivalent architectures in energy management, information processing, and topology, and explicitly calls for benchmarking neuromorphic implementations against conventional architectures to validate predicted gains in energy efficiency and robustness, with protocols such as two-dimensional electronic spectroscopy extended to test quantum coherence in biological computation [9]. Our cost model is exactly the kind of quantitative prediction such benchmarking would need: it states what "efficiency and robustness" should mean for a symmetry-aligned substrate and by how much it should win. The same framework's reflexive-domain construction, building a domain D_∞ as an inverse limit with a fixed-point combinator Y [9], illustrates the general point that computational universality can be achieved inside structurally constrained mathematical universes rather than by unconstrained discrete manipulation.

Finally, the reflection-closure study demonstrates the methodology this paper adopts — prove a classification by explicit small-group arithmetic with every input number declared, and state the limitations with sensitivity for each assumption [10]. Its substantive result is also apt: for finite sets of concurrent lines, closure under mutual reflection always yields exactly the dihedral group $\mathrm{I}_2(n)$ with $n = \mathrm{lcm}_{i,j}\, q_{ij}$ and group order $2n$, nothing else is reachable, and on a rational grid of denominator N the emergence frequencies are exactly $\varphi(q)/N$ for $q \mid N$ [10]. Its framing observation — that model builders in physics and machine learning must choose a symmetry group before fitting begins, that a wrong or hidden choice degrades predictions, and that the choice is normally an unchecked input [10] — is the design gap this paper attacks from the hardware side. A companion epistemological point, that representational models compress and that confusing the map with the territory produces categorical errors [11], cautions the interpretation of our own model: the Noether cost model is a map of substrate behavior, and its assumptions (listed in Section 6) mark where the compression sits.

3. Methods

This paper is analytical. We perform no experiment, no hardware measurement, and no simulation; every number below is either derived here with shown arithmetic or labeled as an assumption feeding a projection. The method has four parts.

Model of the two implementations. We consider a continuous computational task defined by an autonomous system of ODEs $\dot{z} = F(z)$ on a phase space of dimension $2n$, possessing r functionally independent conserved quantities $Q_1, \ldots, Q_r$ with $\nabla Q_1, \ldots, \nabla Q_r$ linearly independent on the working region. Implementation (i), the *digital emulator*, integrates with step size h using a standard (non-symplectic, non-projective) integrator in fixed-precision arithmetic with b bits of mantissa; we assume one effective rounding event per step per conserved quantity, of relative size at most the unit roundoff $u = 2^{-b}$. Implementation (ii), the *symmetry-equivariant processor*, is a physical system whose Lagrangian L is invariant under a group G whose infinitesimal generators correspond to the conserved quantities Q_a; its trajectory is the computation. We assume the processor's only conservation violation is linear physical dissipation with rate constant κ (per unit conserved quantity, relative), so that $\mathrm{d}Q_a/\mathrm{d}t = -\kappa Q_a$.

Choice of cost model. The primary quantitative framework of this paper is the drift-based model: the advantage of symmetry instantiation is measured in conservation-law drift and in the induced error-state geometry, not in a reduction of the Landauer erasure energy. We take this convention because the drift model's inputs (word length, step count, resonator quality factor) are designable hardware parameters with a standard relation $\kappa = \omega/Q_{\mathrm{res}}$, whereas a Landauer-style reduction would require an additional, unestablished assumption that erasure cost scales with the number of constrained degrees of freedom. A Landauer-style reduction factor is therefore treated only as an explicitly labeled hypothesis in Section 6, not as a result.

Derivation. The drift theorem (Section 4.1) is proved by summing the per-step rounding perturbation along the emulator's step sequence and comparing with the Noether identity on the processor side. The Noether side uses the standard identification of conservation-law multipliers with variational symmetries for Lagrangian systems [1]: if L is G-invariant, the multiplier of each Q_a is a symmetry, and the associated continuity equation holds identically on hardware trajectories, so the only drift is the explicitly modeled dissipation. The crossover (Section 4.3) equates the two drift expressions. The volume contraction (Section 4.4) counts discretized error states on and off the conserved manifold.

Assumptions declared up front. (A1) One effective rounding event per step per conserved quantity, of relative magnitude $\leq u$; this is a worst-case accounting, and we also compute the unbiased-rounding random-walk case. (A2) The equivariant processor's dissipation is linear with relative rate κ. (A3) For the numerical projections we assume the parameter set: mantissa $b = 32$ bits; step count $N = 10^6$; resonant frequency $\omega = 10^9\ \mathrm{rad\,s^{-1}}$; resonator quality factor $Q_{\mathrm{res}} = 10^6$; computation wall-clock time $T = 10^{-6}$ s. These are stated assumptions standing in for hardware data, not measurements. (A4) The emulated system's conserved quantities have non-zero gradient on the working region, so that a rounding perturbation of the state produces a first-order perturbation of each Q_a.

Scope. The theorem concerns conservation-quantity drift and the error-state geometry it induces. It does not claim bounds on error growth *within* the conserved manifold, which is governed by the system's Lyapunov exponents and is discussed as a limitation in Section 6.

4. Analysis

4.1 Unit roundoff and per-step drift of the emulator

Input numbers. Mantissa width $b = 32$ bits (assumption A3). Unit roundoff for round-to-nearest:

$$u = 2^{-b} = 2^{-32} = \frac{1}{4\,294\,967\,296}.$$

Computing the decimal value:

$$2^{32} = 4\,294\,967\,296, \qquad u = \frac{1}{4\,294\,967\,296} \approx 2.328306 \times 10^{-10}.$$

Per-step drift. Let Q be a conserved quantity of the emulated system with $|Q|$ the working scale, and let the state be perturbed by a relative rounding of size $\leq u$ at each step. To first order, the induced relative perturbation of Q is bounded by the product of the state perturbation and the condition number of Q on the working region; under assumption (A4) and taking the condition factor as $C \geq 1$ (we set $C = 1$ for the headline numbers, i.e., we assume the conserved quantity is well-conditioned, $\|\nabla Q\|\,\|z\|/|Q| = O(1)$), the per-step relative drift satisfies

$$\left|\frac{\Delta Q_{\text{step}}}{Q}\right| \leq C\,u = 2.328306 \times 10^{-10} \quad (C = 1).$$

Worst-case accumulation. Rounding errors can align in sign, so the worst case after N steps is linear:

$$\left|\frac{\Delta Q_{\text{worst}}(N)}{Q}\right| \leq N\,C\,u.$$

With $N = 10^6$ (assumption A3) and $C = 1$:

$$N\,u = 10^6 \times 2.328306 \times 10^{-10} = 2.328306 \times 10^{-4}.$$

So the worst-case relative conservation drift of a 32-bit emulator after one million steps is 2.33×10^{-4}.

Random-walk accumulation. If rounding errors are unbiased and independent, the drift performs a random walk with root-mean-square

$$\left|\frac{\Delta Q_{\text{rms}}(N)}{Q}\right| = C\,u\,\sqrt{N} = 2.328306 \times 10^{-10} \times \sqrt{10^6} = 2.328306 \times 10^{-10} \times 10^3 = 2.328306 \times 10^{-7}.$$

The gap between 2.33×10^{-4} and 2.33×10^{-7} — three orders of magnitude — is the price of not controlling error sign correlations; a symmetry-exact substrate removes the question entirely, as follows.

4.2 Noether drift of the equivariant processor

Noether side. Let the processor's Lagrangian L be invariant under G with infinitesimal generators whose associated multipliers reproduce $Q_1, \ldots, Q_r$. By the Noether identification of variational symmetries with conservation-law multipliers for Lagrangian systems [1], each Q_a satisfies a continuity equation identically on all trajectories:

$$\frac{\mathrm{d}Q_a}{\mathrm{d}t} = 0 \quad (\text{exact dynamics}).$$

This identity is structural, not numerical: it holds for every trajectory the hardware executes, at every operation count, with no dependence on how many elementary dynamical events the computation comprises. With dissipation (assumption A2), $\mathrm{d}Q_a/\mathrm{d}t = -\kappa Q_a$, giving over a wall-clock time T:

$$\frac{Q_a(T)}{Q_a(0)} = e^{-\kappa T}, \qquad \left|\frac{\Delta Q_{\text{sym}}(T)}{Q}\right| = 1 - e^{-\kappa T} \approx \kappa T \quad (\kappa T \ll 1).$$

Key structural contrast. The emulator's drift is a function of N (operation count); the equivariant processor's drift is a function of T (wall-clock time) and hardware quality only. These are different variables: $N = T/h$ grows as the step size shrinks, while κT does not depend on discretization at all.

4.3 Crossover operation count

Equate the two drift expressions with $C = 1$ and $\kappa T \ll 1$:

$$N^* u = \kappa T \quad \Longrightarrow \quad N^* = \frac{\kappa T}{u}.$$

For a resonant substrate the linear dissipation rate is set by the quality factor, $\kappa = \omega/Q_{\text{res}}$ (standard resonator relation: energy decays as $e^{-\omega t/Q_{\text{res}}}$). Then

$$N^* = \frac{\omega T}{Q_{\text{res}}\, u}.$$

Numerical evaluation (assumptions A3: $\omega = 10^9$ rad s^{-1}, $Q_{\text{res}} = 10^6$, $T = 10^{-6}$ s, $b = 32$):

Step 1 — analog drift: $\kappa T = \dfrac{\omega T}{Q_{\text{res}}} = \dfrac{10^9 \times 10^{-6}}{10^6} = \dfrac{10^3}{10^6} = 10^{-3}$.

Step 2 — crossover: $N^* = \dfrac{10^{-3}}{2.328306 \times 10^{-10}} = \dfrac{10^{-3} \times 4\,294\,967\,296}{1} = 4\,294\,967.296 \approx 4.295 \times 10^6$.

(Computation: $1/u = 2^{32} = 4\,294\,967\,296$, so $N^* = 10^{-3} \times 4\,294\,967\,296 = 4.294967296 \times 10^6$.)

Rounding to three significant figures, $N^* \approx 4.30 \times 10^6$ steps. Beyond roughly 4.3 million steps at the assumed parameters, the symmetry-equivariant substrate has strictly lower conservation drift than the 32-bit emulator; below it, the emulator wins on this metric. The crossover scales linearly with Q_{res} (better resonators favor the analog side later) and inversely with b (more mantissa bits favor the emulator later): $N^* = 2^b\, \omega T / Q_{\text{res}}$.

4.4 Error-state volume contraction

Each functionally independent conserved quantity confines the state to a codimension-one manifold, removing one independent error direction. Count discretized error states: if the unconstrained error ball of relative radius u has volume V_0 in the $2n$-dimensional error space, then with r conserved quantities the reachable error volume is the intersection with an $(2n - r)$-dimensional slice of thickness u per constraint:

$$V_r = V_0\, u^r, \qquad \frac{V_r}{V_0} = u^r.$$

For $r = 1$ (assumption A3, $b = 32$):

$$\frac{V_1}{V_0} = u = 2.328306 \times 10^{-10}.$$

For $r = 2$:

$$\frac{V_2}{V_0} = u^2 = \left(2.328306 \times 10^{-10}\right)^2 = 5.421010 \times 10^{-20}.$$

(Computation: $2.328306^2 = 5.42101$, and $\left(10^{-10}\right)^2 = 10^{-20}$.)

Two independent conserved quantities contract the reachable error volume by twenty orders of magnitude relative to the unconstrained case. This contraction is the geometric content of the claim that conserved quantities act as native error control: the emulator must buy the same reduction with bits or with projection steps, at per-operation cost.

4.5 Sensitivity to the assumed parameters

The crossover is $N^* = 2^b\, \omega T / Q_{\text{res}}$. Holding the other inputs fixed:

- Doubling the mantissa, $b = 64$: $N^* = 2^{64} \times 10^{-3} = 1.8446744 \times 10^{19} \times 10^{-3} = 1.845 \times 10^{16}$ steps. Double precision pushes the crossover far beyond any plausible run, so the analog advantage under this metric applies to moderate-precision, very-long-run emulation.
- A ten-times better resonator, $Q_{\text{res}} = 10^7$: $N^* = 10^{-4}/u = 4.295 \times 10^5$ steps; the analog side wins ten times sooner because its drift is ten times smaller.

- A ten-times longer computation, $T = 10^{-5}$ s: $\kappa T = 10^{-2}$, $N^* = 4.295 \times 10^7$ steps; the crossover rises with wall-clock duration because the analog drift budget κT grows.

The worst-case drift claim Nu is sensitive to the condition factor C: if the conserved quantity is ill-conditioned on the working region ($C \gg 1$), both drift bounds scale up by C and the crossover falls by C; the qualitative conclusion (zero structural drift versus linear numerical drift) is unaffected, but the numerical crossover is not meaningful without a conditioning audit of the specific Q.

5. Results

The quantities derived in Section 4, with their status, are:

Quantity	Symbol	Value	Status
Unit roundoff, $b = 32$	u	2.328306×10^{-10}	Derived (Section 4.1)
Worst-case emulator drift, $N = 10^6$	Nu	2.328306×10^{-4}	Projection (assumes A1, A3)
Random-walk emulator drift, $N = 10^6$	$u\sqrt{N}$	2.328306×10^{-7}	Projection (assumes unbiased rounding, A3)
Analog drift over $T = 10^{-6}$ s	κT	1.0×10^{-3}	Projection (assumes A2, A3)
Crossover operation count	N^*	4.295×10^6	Projection (assumes A1–A3)
Error-volume contraction, $r = 1$	u^1	2.328306×10^{-10}	Derived given $b = 32$
Error-volume contraction, $r = 2$	u^2	5.421010×10^{-20}	Derived given $b = 32$
Crossover at $b = 64$	N^*	1.845×10^{16}	Projection (assumes A2–A3, $b = 64$)

Three results carry the paper. First, the structural result: a G-equivariant processor's conservation drift is exactly zero on exact dynamics and κT with dissipation, independent of operation count, while the emulator's drift grows at least as $u\sqrt{N}$ and at most as Nu. Second, the crossover: at the assumed parameters the symmetry-aligned substrate wins on conservation drift beyond $N^* \approx 4.30 \times 10^6$ steps, and the crossover is tunable through b, Q_{res}, and T in known directions. Third, the geometric result: each conserved quantity contracts the reachable error volume by u, so even one enforced invariant is worth roughly ten orders of magnitude of error-state exclusion at 32-bit resolution, and two are worth twenty.

These results demonstrate that symmetry choice is a quantitative error-control design parameter, not a qualitative preference: the model states which substrate wins, at what operation count, and by how much, under declared assumptions.

6. Discussion

What the model does not claim. The drift theorem bounds conservation-quantity drift only. Error growth *within* the conserved manifold — the directions the symmetry does not remove — is governed by the system's Lyapunov exponents, and a chaotic emulator and a chaotic equivariant processor both amplify in-manifold perturbations; the model buys constraint enforcement, not chaos suppression. The comparison is also metric-specific: it measures conservation drift, not solution accuracy in general, and a task that does not value any conserved quantity gains nothing from the model.

Falsifiability. The main result would be falsified by any of three observations. First, if a physical G-equivariant processor were built whose conserved quantity drifts linearly with the number of internal dynamical events rather than with wall-clock time and quality factor, the structural contrast in Section 4.2 would be wrong — this would mean the hardware Lagrangian is not actually G-invariant, or that the conserved quantity is not among the Noether multipliers, and the failure would be diagnosable by recomputing the multiplier identification [1]. Second, if the measured drift of a long-run digital emulator were sub-linear in N in worst-case regimes (not merely the unbiased-rounding random walk

$u\sqrt{N}$, which the model already accommodates), the per-step rounding model (A1) would overstate the emulator's cost and the crossover N^* would move; the model survives as long as drift is at least $u\sqrt{N}$, and dies if drift is $O(u)$ independent of N. Third, if the dissipation of a candidate substrate were not linear in Q_a — for example, if drift grew as $\sqrt{t}$ or t^2 through nonlinear loss channels — the crossover formula $N^* = 2^b\,\omega T/Q_{\text{res}}$ would need replacement, though the qualitative statement (analog drift bounded by a hardware constant, emulator drift unbounded in N) would survive any bounded-dissipation model.

A labeled hypothesis: Landauer-style energy reduction. The research program this paper belongs to conjectures an energy-dissipation advantage for symmetry-equivariant processors. Our drift model does not establish one. A Landauer-style hypothesis would be: if erasing an error state costs $k_{\text{B}}T_{\text{env}}\ln 2$ per bit of removed uncertainty, and r conserved quantities remove a fraction $1-u^r$ of the error-state volume, then the per-operation erasure energy of the equivariant processor is lower by a factor related to u^r. We state this as a hypothesis only: it requires an additional assumption — that erasure cost scales with the constrained volume — that we have not derived and that no cited work establishes. It is not a result of this paper.

Limitations. The result rests on the following assumptions, with sensitivity noted for each:

- *Rounding model (A1).* One effective rounding event per step per conserved quantity, of relative size $\le u$. If a step involves m effective roundings, the worst-case bound becomes Nmu and the crossover falls by m; the qualitative conclusion is unchanged, the numerical crossover is not meaningful without an operation-level audit of the integrator.
- *Conditioning (A4, factor C).* Both drift bounds scale linearly with C; an ill-conditioned conserved quantity ($C \gg 1$) can erase the numerical advantage entirely at moderate N. The crossover is meaningless without a conditioning audit of the specific Q_a.
- *Linear dissipation (A2).* The processor side of the comparison is exactly κT only if dissipation is linear with relative rate $\kappa = \omega/Q_{\text{res}}$. Nonlinear or state-dependent loss channels invalidate the closed form.
- *Assumed parameter set (A3).* $b = 32$, $N = 10^6$, $\omega = 10^9\ \text{rad}\,\text{s}^{-1}$, $Q_{\text{res}} = 10^6$, $T = 10^{-6}$ s are stated assumptions, not measurements; every headline number inherits their uncertainty. The crossover is most sensitive to b (exponentially, $N^* = 2^b\,\omega T/Q_{\text{res}}$) and linearly sensitive to Q_{res} and T.
- *Independence of conserved quantities.* The volume contraction u^r counts only functionally independent invariants with linearly independent gradients; dependent invariants buy nothing additional.
- *Metric scope.* The comparison measures conservation drift and error-state volume only, not wall-clock throughput, energy per operation, or general solution accuracy; a substrate could win on drift and lose on all other metrics.
- *Map-territory caution.* The model is a compression of substrate behavior, and confusing the model with the substrate would repeat the categorical error the map-territory literature warns against [11]; the assumptions above mark where the compression sits.

Open questions. Three remain. (i) Does the crossover survive when the emulator uses a structure-preserving integrator that conserves Q_a to machine precision by construction? Such an integrator moves the emulator's drift to the projection steps it performs, so the comparison becomes one of projection frequency versus dissipation rate; we have not modeled it. (ii) What is the correct cost model when the emulated system's symmetry group differs from the substrate's? A mismatch means the substrate conserves the wrong quantities, and the drift comparison inverts. (iii) Does the u^r volume contraction translate into a fault-tolerance statement, as the topological-protection conjecture in the source program suggests? The contraction is a static geometric count; converting it to a dynamical error-correction claim requires a noise model, which we have not supplied.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The users are designers of analog and neuromorphic accelerators — photonic resonator networks, oscillator-based computing systems, and symplectic analog substrates — and, secondarily, numerical analysts choosing among structure-preserving integrators. If the model holds, the design change is concrete: the symmetry group of the substrate becomes a declared design parameter, audited at design time the way word length is audited in digital design, because the model states which conserved quantities the substrate enforces for free and at what operation count that enforcement beats a digital emulator of given mantissa width. The decision it changes is the choice between implementing a long-running continuous simulation on a digital emulator versus a symmetry-aligned analog substrate: at the assumed parameters the analog side wins on conservation drift beyond $N^* \approx 4.30 \times 10^6$ steps, so a workload profile with more than roughly 4.3 million effective steps per conserved quantity favors the analog route. The benchmarking call in the categorical-equivalence framework — that neuromorphic implementations should be benchmarked against conventional architectures to validate predicted gains in energy efficiency and robustness [9] — needs exactly the kind of quantitative prediction this model supplies: it names the metric (conservation drift), the crossover, and the scaling.

(b) Current readiness. This work is purely analytical and has not been peer reviewed: by the NASA Technology Readiness Level definitions it sits at TRL 1–2 — basic principles observed and reported, and concept formulated. No hardware has been built, measured, or simulated under this model; analytical work of this kind does not exceed TRL 3, and nothing here approaches it. The resonator relations used ($\kappa = \omega/Q_{\text{res}}$) are standard, but their assembly into the crossover model is unvalidated against any physical processor.

(c) The smallest next test. The smallest de-risking experiment is a single-substrate drift measurement: take one photonic or electronic resonator network whose Lagrangian has a known continuous symmetry (for example a U(1)-symmetric oscillator pair conserving a norm-like quantity), run it for a wall-clock time T with measured Q_{res} and ω, and measure the relative drift of the conserved quantity; in parallel, integrate the same system on a 32-bit digital emulator for $N = 10^6$ steps and measure its drift. The model is refuted if the analog drift exceeds $\kappa T = \omega T/Q_{\text{res}}$ by more than the conditioning factor C, or if the emulator's worst-case drift falls below $u\sqrt{N}$; it is supported if both bounds hold within the declared C. A study would also measure the condition number $\|\nabla Q\|\,\|z\|/|Q|$ on the working region, since Section 4.5 shows the crossover is not meaningful without it.

(d) Binding constraints. In the order they would likely bind: (i) *conditioning audit* — the conserved quantity must be well-conditioned on the working region, or the drift bounds inflate by C (Section 4.5); (ii) *loss-channel linearity* — the substrate must exhibit linear dissipation at the operating amplitude, since nonlinear loss breaks the closed-form crossover; (iii) *symmetry fidelity* — fabrication tolerances in resonator networks break exact G-invariance, and the model says nothing about how much symmetry breaking re-introduces drift; this is an open question the grounding does not quantify; (iv) *I/O conversion* — reading the analog result back into digital form reintroduces rounding at the boundary, an error channel outside the model; (v) *cost and regulation* — the grounding supplies no cost, energy, or regulatory data for analog substrates, so any deployment economics is an open question, not a claim.

(e) Order-of-magnitude performance bounds (projections). Under the assumptions of Section 4 (A1–A3), the model projects: conservation drift of the analog substrate at or below 1.0×10^{-3} relative over $T = 10^{-6}$ s, independent of operation count; emulator drift between 2.33×10^{-7} (random walk) and 2.33×10^{-4} (worst case) at $N = 10^6$ steps; and a crossover at $N^* = 2^b\,\omega T/Q_{\text{res}}$, ranging from 4.295×10^5 steps ($Q_{\text{res}} = 10^7$) to 1.845×10^{16} steps ($b = 64$). These are projections from the stated assumptions with no uncertainty quantification beyond the sensitivity analysis in Section 4.5; they are not measurements.

(f) Risks and misuse paths. The main misuse is reading the crossover as a general performance claim: the model measures conservation drift only, and quoting $N^* \approx 4.30 \times 10^6$ as "analog beats digital" without the drift metric, the conditioning audit, and the declared assumptions would overstate the result. A second risk is symmetry-washing: declaring a substrate "symmetry-aligned" without verifying that the hardware Lagrangian is actually G-invariant and that the claimed Q_a are among its Noether multipliers [1] — the same unchecked-input failure mode the reflection-closure study identifies for model builders who choose a symmetry group without checking it [10].

8. Conclusion

This paper derived a Noether-based cost model comparing a b-bit digital emulator of a conservative continuous system against a G-equivariant physical processor that instantiates the same conservation laws. The emulator's conserved-quantity drift is bounded by Nu in the worst case and $u\sqrt{N}$ under unbiased rounding, giving 2.33×10^{-4} and 2.33×10^{-7} respectively at $N = 10^6$ steps with $b = 32$; the equivariant processor's drift is zero on exact dynamics and $\kappa T = \omega T / Q_{\text{res}} = 10^{-3}$ under linear dissipation, independent of operation count. Equating the two gives a closed-form crossover $N^* = 2^b \, \omega T / Q_{\text{res}} \approx 4.30 \times 10^6$ steps at the assumed parameters, tunable exponentially through b and linearly through Q_{res} and T. Each functionally independent conserved quantity contracts the reachable error-state volume by u per invariant — ten orders of magnitude at $r = 1$ and twenty at $r = 2$ for a 32-bit scale.

The conclusion these numbers support is narrow but designable: symmetry choice is a quantitative error-control parameter for physical processors, on the same footing as word length, with a computable crossover against any specified digital baseline. The result is bounded by its assumptions — one rounding event per step, well-conditioned invariants, linear dissipation, and a parameter set that is declared rather than measured — and it measures conservation drift, not general accuracy or energy. The falsifiable core is the structural contrast of Section 4.2: analog drift scales with wall-clock time and quality factor, emulator drift with operation count. A single-substrate drift measurement against a matched digital integration, as specified in Section 7, would confirm or refute it. The analysis may also serve adjacent fields where conserved quantities discipline long computations — long-horizon geophysical and molecular simulation and analog signal processing — insofar as their invariants fall in the volume, surface, and curve classes classified for three-dimensional PDE systems [2].

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 10 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] Conservation laws of scaling-invariant field equations. arXiv:math-ph/0303066v3. https://arxiv.org/abs/math-ph/0303066v3 [2] On the different types of global and local conservation laws for partial differential equations in three spatial dimensions: review and recent developments. arXiv:1803.08859v2. https://arxiv.org/abs/1803.08859v2 [3] Conserved norms and related conservation laws for multi-peakon equations. arXiv:1707.05389v2. https://arxiv.org/abs/1707.05389v2 [4] Darboux Transformations, Infinitesimal Symmetries and Conservation Laws for Nonlocal Two-Dimensional Toda Lattice. arXiv:nlin/0112035v3. https://arxiv.org/abs/nlin/0112035v3 [5] Group foliations, invariant solutions, and conservation laws of the geopotential forecast equation. arXiv:2310.14562v3. https://arxiv.org/abs/2310.14562v3 [6] Wave Function Collapse, Correlating Interactions, and Conservation Laws. arXiv:2102.11370v4. https://arxiv.org/abs/2102.11370v4 [7] Sub-symmetries II. Sub-symmetries and Conservation Laws. arXiv:1705.02289v1. https://arxiv.org/abs/1705.02289v1 [8] https://papers.qnfo.org/papers/2-context. QNFO: 2. Historical Context and Literature Review. [9] https://papers.qnfo.org/papers/biological-and-engineered-computational-systems. QNFO: Biological and Engineered Computational Systems. [10] https://papers.qnfo.org/papers/reflection-closure-of-concurrent-line-sets-exact-dihedral-emergence-grid-frequen. QNFO: Reflection Closure of Concurrent Line Sets: Exact Dihedral Emergence, Grid Frequency Laws, and the Finite/Infinite Dichotomy. [11] https://papers.qnfo.org/papers/principia-ontologica. QNFO: Principia Ontologica.