Defensive Publication DP-2026-02
CC0 1.0 · public-domain dedication
7 July 2026

DEFENSIVE PUBLICATION · ENERGY-NATIVE COMPUTE

A Thermodynamic Bound on Reproducibility in Energy-Native Computation, and the Crossings-Tax Spectrum

Collapse the crossings between source and compute state and efficiency rises but noise floods in. Reproducibility is bought back not with a conversion crossing but with a dissipation-built basin, at a price with a lower bound.

David Jean Charlot, PhD · Michael Hiner

The Charlot Lab & The Hiner Lab, Institute for Physical AI @ BMI

Correspondence: contact@physicalai-bmi.org · physicalai-bmi.org · Interactive companion: physicalai-bmi.org/research/hiner-lab#topic-energy

Energy-delivery-to-compute schemes are ordered by one axis: the number of domain crossings that survive between the energy source and the computational state variable, each crossing a conversion loss and a reproducibility checkpoint. Collapsing crossings raises efficiency and, along the same diagonal, forfeits the interfaces at which state could have been re-clocked or restored — so the maximally efficient substrate (identity, zero crossings) inherits the raw noise of its carrier. The disclosed object is the resolution and its price: define the committed observable as the label of the attractor basin a dissipative substrate settles into, not the microstate; then supply perturbations below the barrier are physically rejected, the settling is the error correction, and attestation binds only the label at a cost of order the decision entropy. The barrier height that rejects supply noise of a given intensity is a thermodynamic lower bound on the energy per reproducible commit — a Landauer-analog distinct from the erasure bound — which forces two consequences: reproducibility and reversibility are opposed (a non-dissipative substrate builds no basins), and free invariance is paid in representational density. All constituent physics is public; no term is enclosable.

1. The crossings-tax axis

Let a computing system map an energy source to a computational state variable through a chain of physical domains. Write K for the number of domain crossings that survive between the two. Each crossing c carries a conversion inefficiency $\lambda_c \in (0, 1)$ and freezes an interface, so the end-to-end efficiency is multiplicative,

$$\eta = \prod_{c=1}^{K} (1 - \lambda_c),$$

and is monotonically improved by deleting crossings. The spectrum orders mechanisms by crossings remaining, high to low: **S0** separated (source · supply · regulated rail · digital compute); **S1** harvested-separated; **S2** conditioning-collapsed (bursty supply drives asynchronous logic, no clean rail); **S3** storage-collapsed (the store is the state); **S4** co-integrated (source and compute share a substrate, distinct mechanisms); **S5** identity (the source variable *is* the compute-state variable, $K = 0$). S5 is the asymptote and the object of interest.

2. The carrier × rung matrix

Crossed with the physical carrier, the spectrum has forbidden and empty cells that locate the unbuilt work. The S5-identity mechanism and its status per carrier:

CARRIER	S5-IDENTITY MECHANISM (SOURCE VAR = COMPUTE VAR)	STATUS
Photonic	incident light performs the linear transform on its own energy	demonstrated (diffractive optical networks)
Electrochem.	redox / ion gradient powers and computes	demonstrated in biology (neurons)
Ionic	Soret-driven ion separation *is* the memristive state	empty — implementable (both halves published)
Thermal	gradient powers and computes in-domain	empty — nascent
Mechanical	harvested strain performs the logic	near (mechano-ionic / phononic)
Spin / magnetic	spin texture stores energy and computes	frontier (magnonic)
RF / EM	incident field powers and modulates the computation	partial (backscatter)
Nuclear	—	**forbidden** — beta decay is power-only; caps at S4

The empty-but-implementable cells — ionic and thermal — are the flagged prior art: in each, two literatures name one mechanism twice (harvesting versus computing) and no published device makes one coordinate serve both. Nuclear is the only carrier forbidden to identity, which is precisely why a betavoltaic constant-power floor is the deterministic reference: pure power, no compute coupling, no noise to inject.

3. The diagonal (design law)

The spectrum has a diagonal. Every crossing deleted to remove the conversion tax is a crossing at which the state could have been re-clocked, re-quantized, or restored to enforce reproducibility. Efficiency and bit-invariance are therefore bought with the same coin, the retained interface: S0 is maximally reproducible and maximally wasteful; S5 is maximally efficient and inherits the raw physics — including the noise — of its carrier. The naive repair, converting and rounding at an added boundary, reintroduces exactly the crossing that was deleted. The remaining sections disclose an invariance that costs no crossing.

4. Basin-label invariance

Let the substrate be dissipative, descending a free-energy landscape V with discrete attractor basins $\{B_\ell\}$. Define the committed observable as the basin label rather than the microstate,

$$\ell = \Phi(\mathbf{x}) = \{\, \ell : \mathbf{x} \in B_\ell \,\}, \qquad H = \text{SHA-256}\big(\ell(t)\big) \stackrel{!}{=} H^{\text{ref}}.$$

Three mechanisms make the label invariant at no domain crossing, because each is intrinsic to the substrate. **(i) Basin coarse-graining:** a perturbation below the barrier is rejected as the dynamics roll back down; the settling is the error correction, and the landscape does the rounding that a converter would otherwise do. **(ii) Physical redundancy:** many noisy elements summed on a shared electrode average by the law of large numbers, and the summation is a Kirchhoff junction — a wire — so independent noise never moves the committed label; only a common-mode excursion can. **(iii) Hysteresis as free quantization:** a hysteretic device is a physical Schmitt trigger, rejecting noise inside the loop width without an analog-to-digital conversion. Attestation then pays exactly one crossing, and it is asymptotically free: signing the label costs $O(\log |\{B_\ell\}|)$ bits of decision entropy, not the joules-heavy analog trajectory.

5. The bound (a Landauer-analog for reproducibility)

Escape from a basin under supply noise of intensity D_σ (the noise variance density, an effective temperature) follows Kramers' rate[4],

$$r = \frac{\omega_0\,\omega_b}{2\pi\gamma}\, e^{-\Delta E_b / D_\sigma}.$$

Holding the label over a commit window τ with failure probability at most ε requires $1 - e^{-r\tau} \le \varepsilon$, hence a minimum barrier, and — since each rejected perturbation dissipates at least the barrier it climbs back down — a minimum energy per invariant commit:

$$\boxed{E_{\text{inv}} \;\ge\; \Delta E_b^{\min} \;=\; D_\sigma \,\ln \frac{\omega_0 \omega_b\,\tau}{2\pi\gamma\,\varepsilon} \;\xrightarrow[D_\sigma \to k_B T]{}\; \Theta\big(k_B T \ln(1/\varepsilon)\big).}$$

This is a thermodynamic lower bound on the energy required to make an output invariant against supply noise of a given amplitude, and it is distinct from Landauer's $k_B T \ln 2$ per erased bit[1]: Landauer prices erasure, this prices reproducibility. Spend less than $\Delta E_b^{\min}$ and the basins are too shallow to reject the noise — the invariance is not available to purchase at any lower price. Whether this bound is tight, or can be beaten by a coarse-graining cleverer than $k_B T \ln(\text{states})$, is stated as the open theoretical question, not asserted as settled.

6. The reversibility inversion

The bound inverts a premise the low-power field treats as settled. Reversible and adiabatic computing minimize energy per operation by not dissipating[2]; but a non-dissipative substrate has no attractor basins, and so preserves supply noise instead of coarse-graining it. Dissipation is what builds the basins that reject the noise for free. The substrate that buys boundary-invariance without a crossing is therefore minimally dissipative but deliberately *not* reversible, sitting above the Landauer floor by exactly the barrier height it needs,

$$Q_{\text{commit}} \;\gtrsim\; \Delta E_b^{\min} \;=\; D_\sigma \ln(1/\varepsilon) \;+\; \text{const},$$

so that reproducibility and reversibility pull against each other at S5. The attractor-network lineage of this construction is classical[3]; what is disclosed is its use as a crossing-free invariance primitive priced by the bound above.

7. Representational density (the honest cost)

Free invariance is paid in bits. Deep, well-separated basins occupy the state range at a coarser spacing $w(\Delta E_b)$, increasing in barrier height, so the number of distinguishable states and hence the information capacity per device fall as the barrier rises,

$$N_{\text{states}} = \frac{R}{w(\Delta E_b)}, \qquad \text{bits/device} = \log_2 N_{\text{states}} \downarrow \;\text{ as } \Delta E_b \uparrow .$$

This is the true Pareto surface of energy-native computation: representational density traded against noise margin, a landscape-geometry tradeoff in which barrier height replaces the guardband of a regulated design.

8. Claim

The disclosed object is the following, taken together and priced thermodynamically rather than in raw energy: (a) the crossings-tax ordering S0–S5 and the carrier matrix that locates the empty-but-implementable identity cells; (b) the basin-label invariance construct, in which the committed and attested observable is the attractor-basin label of a dissipative substrate and the three intrinsic mechanisms of §4 supply reproducibility at no domain crossing; and (c) the bound $E_{\text{inv}} \geq D_\sigma \ln(1/\varepsilon)$ with its two corollaries — the reversibility inversion and the density–margin Pareto. All constituent relations — the multiplicative efficiency, Kramers' rate, the Landauer floor, the Kirchhoff sum, the Nernst/Soret and photonic identity mechanisms — are public physics; no term is enclosable.

References

1. R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* 5(3), 183–191, 1961. doi:10.1147/rd.53.0183.
2. C. H. Bennett. Logical reversibility of computation. *IBM J. Res. Dev.* 17(6), 525–532, 1973. doi:10.1147/rd.176.0525.

3. J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *PNAS* 79(8), 2554–2558, 1982. doi:10.1073/pnas.79.8.2554.
4. H. A. Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica* 7(4), 284–304, 1940. doi:10.1016/S0031-8914(40)90098-2.
5. X. Lin, Y. Rivenson, N. T. Yardimci, M. Veli, Y. Luo, M. Jarrahi, A. Ozcan. All-optical machine learning using diffractive deep neural networks. *Science* 361, 1004–1008, 2018. doi:10.1126/science.aat8084. (S5 photonic identity.)
6. T. M. Kamsma, W. Q. Boon, T. ter Rele, C. Spitoni, R. van Roij. Iontronic neuromorphic signaling with conical microfluidic memristors. *Phys. Rev. Lett.* 130, 268401, 2023. arXiv:2301.06158.
7. Nanofluidic logic with mechano–ionic memristive switches. *Nature Electronics* 7, 271–278, 2024. doi:10.1038/s41928-024-01137-9.
8. Anionic entanglement-induced giant thermopower in ionic thermoelectric materials. *eScience* 3, 100169, 2023. doi:10.1016/j.esci.2023.100169.
9. 14C diamond as an energy-converting material in a betavoltaic battery: a first-principles study. *AIP Advances* 13, 115314, 2023. doi:10.1063/5.0177302.

AI-use disclosure. Preparation of this disclosure used a large language model (Claude, Anthropic) for drafting and editing text, formalizing the argument, and preparing the interactive companion. Cited references were checked to resolve to their sources. The authors reviewed the content and are solely responsible for it. Consistent with ICMJE, COPE, and IEEE guidance, the model is a tool and is not credited as an author.

Institute for Physical AI @ BMI · The Charlot Lab & The Hiner Lab
503 McKeever Rd, Arcola, TX 77583, USA
physicalai-bmi.org · contact@physicalai-bmi.org