INSTITUTE FOR PHYSICAL AI @ BMI
The Charlot Lab & The Hiner Lab

Technical Report TR-2026-05
Survey / Review · Preprint v1
7 July 2026

ENERGY-NATIVE COMPUTE

# Energy-Native Computation and the Price of a Reproducible Bit: A Survey and Research Position

Order computation by the crossings between its energy and its state; find that the most efficient substrate is the noisiest; and price the reproducibility you must buy back.

**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 · Formal program: Defensive Publication DP-2026-02

**Abstract.** In a low-power substrate, energy conversion and computation are frequently the same physical mechanism, described by two research communities that do not cite each other: what one calls harvesting, the other calls computing. This report organizes the field with a single axis — the number of domain crossings that survive between the energy source and the computational state variable — and surveys the resulting spectrum from fully separated systems to the asymptote of identity, where the source variable is the state variable and no crossing remains. Identity is not speculative: it is demonstrated for light and for biological ion gradients, and it is empty but implementable for the ionic and thermal carriers, where both halves are already published. The spectrum has a diagonal that is a design law: every crossing deleted to remove the conversion tax is a crossing at which state could have been restored to enforce reproducibility, so efficiency and bit-invariance are bought with the same coin and the most efficient substrate inherits the raw noise of its carrier. The report's research position is a way to buy reproducibility back without a crossing — commit the label of the attractor basin a dissipative substrate settles into, not the microstate — and a thermodynamic bound on its price: the barrier height that rejects supply noise of a given amplitude is a Landauer-analog for reproducibility, with two consequences, that reproducibility and reversibility are opposed, and that free invariance is paid in representational density. We are explicit about maturity and about what remains an open theoretical question. The report reports no new experimental measurements; the full program is released separately as a public-domain defensive publication.

## 1. Introduction

An embodied machine spends most of its energy not on arithmetic but on the infrastructure around it: converting a source into a rail, storing charge, moving heat, and conditioning a supply into the clean, constant form that digital logic expects. Each of those steps is a domain crossing, and each is a loss. For a datacenter the losses are a facilities line item; for a machine that carries its own power they are the design. This report takes seriously the possibility that the crossings themselves — not the transistors — are the thing to remove, and it follows that idea to its limit, where the distinction between the energy system and the computer disappears.

The observation that motivates the survey is that this limit is not exotic. A diffractive optical network performs its matrix multiply on the energy of the light passing through it[5]; a neuron computes on the same ion gradient that powers it. In both, the source variable and the computational state variable are one coordinate. Between those existence proofs and today's separated hardware lies a spectrum, and reading the spectrum honestly turns up both an opportunity — carriers where identity is published in halves but never assembled — and a problem the whole in-materia field currently hand-waves: that the efficiency won by collapsing crossings is paid for in reproducibility. The report maps the first and prices the second.

## 2. Scope and method

This is a survey with a stated research position, not an experimental paper. It reviews published work across several energy carriers and the classical statistical mechanics that connects energy delivery to reproducible output[4,1], favors a primary source per claim, and is explicit about the maturity of each capability. The report advances a position in Sections 5–7 and reports no original measurements. The interactive companion referenced throughout is an illustrative model, labeled as such; the formal statement of the spectrum, the invariance construct, and the bound is released as a separate public-domain defensive publication (DP-2026-02) so that it stands as prior art. Its principal limitation is that the central quantitative claim — that the barrier height rejecting a given supply noise is a genuine lower bound on the energy per reproducible commit — is advanced as a well-motivated conjecture with a scaling argument, not a proved theorem; the report says so where it matters.

## 3. The crossings-tax spectrum

Order every energy-to-compute scheme by the number of domain crossings $K$ that survive between the source and the computational state variable. Each crossing carries a conversion inefficiency and freezes an interface, so efficiency is multiplicative in the crossings and improves as they are deleted. The rungs, from most crossings to fewest, are: separated (S0), harvested-separated (S1), conditioning-collapsed (S2, a bursty supply drives asynchronous logic with no clean rail), storage-collapsed (S3, the store is the state), co-integrated (S4, source and compute share a substrate as distinct mechanisms), and identity (S5, the source variable is the state variable). Crossing the spectrum with the physical carrier gives a matrix whose empty and forbidden cells are the actual content (Table 1 and Figure 1).

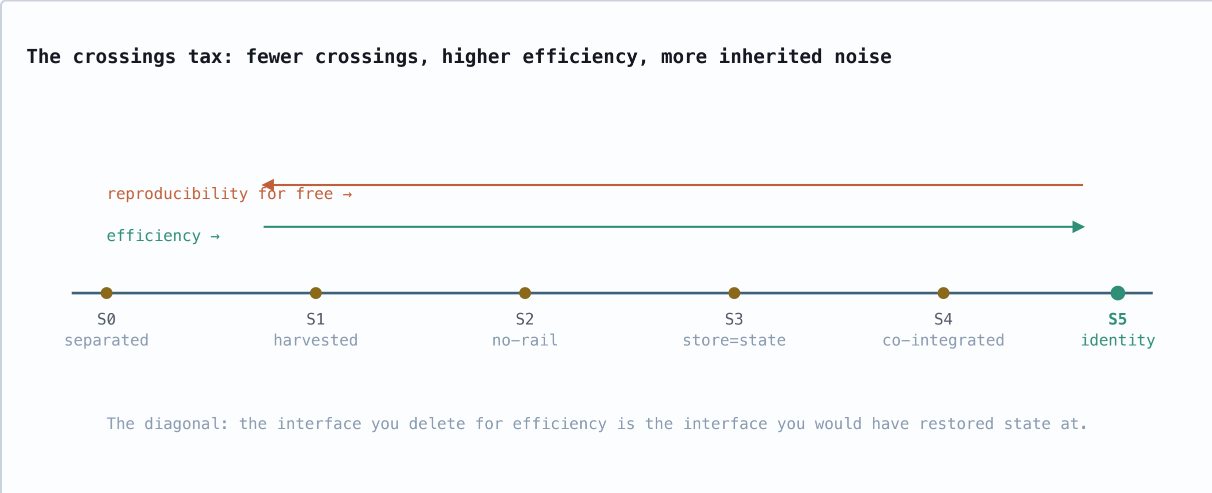


**Figure 1.** The crossings spectrum. Moving toward identity (S5) raises efficiency and strips away the interfaces at which reproducibility could have been enforced, so the substrate inherits its carrier's raw noise. An interactive version that lets the reader trade efficiency against noise and watch the reproducibility floor is at physicalai-bmi.org/research/hiner-lab.

**Table 1.** The S5-identity mechanism per carrier, and its status. The empty-but-implementable cells locate the unbuilt prior art.

| CARRIER | S5-IDENTITY MECHANISM | STATUS |
|---|---|---|
| Photonic | light performs the linear transform on its own energy | demonstrated[5] |
| Electrochemical | ion gradient powers and computes | demonstrated in biology |
| Ionic | Soret-driven separation is the memristive state | empty — implementable[7,6] |
| Mechanical | harvested strain performs the logic | near[8] |
| Thermal | gradient powers and computes in-domain | empty — nascent |
| Spin / magnetic | spin texture stores energy and computes | frontier |
| Nuclear | — | forbidden; caps at S4[9] |

## 4. Where identity is empty but implementable

The sharpest opportunities are the empty cells whose halves are already peer-reviewed. In the ionic carrier, ionic thermoelectrics separate cations and anions along a conductor by the Soret effect, producing a thermopower one to two orders above the electronic Seebeck coefficient[7]; independently, nanofluidic iontronic memristors compute because field-driven ion migration through a channel is itself a memory state, as in the conical microfluidic memristor[6] and the mechano-ionic switch[8]. The driving variable is the same ion in the same channel; one field books its displacement as harvested charge, the other as a synaptic weight, and no published device makes the Soret-driven separation double as the memristive state so that a gradient powers and writes a computation in one coordinate. Because both halves exist, an integrated element is a synthesis rather than a discovery gamble. Nuclear sits at the opposite corner: beta decay is power-only, so a carbon-14 diamond betavoltaic delivers constant micropower for decades with no compute coupling and no noise to inject[9] — which is exactly why it is the natural deterministic reference at the bottom of the spectrum.

## 5. The diagonal: efficiency and reproducibility are one coin

The spectrum's organizing insight is a diagonal. Descend it, and every crossing deleted to remove the conversion tax is also a crossing at which the state could have been re-clocked, re-quantized, or checkpointed to enforce reproducibility. The two properties are bought with the same coin, the retained interface. So the separated system is maximally reproducible and maximally wasteful, and the identity substrate is maximally efficient and inherits the raw physics — including the noise, the device-to-device variation, and the drift — of its carrier. The naive repair, converting and rounding at an added output stage, simply reintroduces the crossing that was deleted to gain the efficiency. The question the position answers is whether reproducibility can be recovered without paying that crossing back.

> The central tension of energy-native computation, stated plainly: the substrate that removes the conversion tax is the substrate that removes the places you would have made the computation reproducible. Efficiency is not free of a reproducibility cost; the cost is just paid in a different currency, and the rest of this report is about that currency.

## 6. Buying invariance back without a crossing

The move is to stop demanding bit-identity of the noisy microstate and to demand invariance only of an equivalence class that the physics produces for free. Let the substrate be dissipative and descend a free-energy landscape into discrete attractor basins. Define the committed observable as the basin label, not the trajectory. Three mechanisms then supply reproducibility at no domain crossing, because each is intrinsic to the substrate. First, basin coarse-graining: a supply perturbation below the barrier is physically rejected as the dynamics roll back down, so the settling is the error correction and the landscape does the rounding a converter would otherwise perform. Second, 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 and only a common-mode excursion can. Third, hysteresis: a hysteretic device is a physical Schmitt trigger that quantizes without an analog-to-digital stage. Attestation then binds only the label, at a cost of order the decision entropy — a few bits — rather than the joules-heavy analog state, so the one crossing it pays is asymptotically free.

This reframes the property that verification actually needs. Invariance at identity is not bit-identical microstate; it is same-basin, same-outcome-class, which is strictly weaker than digital bit-identity and is exactly what a certification step requires, since verification asks only that the checked outcome equal the deployed outcome, never that every intermediate microstate match. Invariance becomes a property of landscape geometry — engineered barrier heights — rather than of a checkpoint operation, and barrier height is the design parameter that replaces the guardband.

## 7. The price of a reproducible bit

The barrier is not free, and its price has a floor. Escape from a basin under supply noise of intensity $D_\sigma$ follows Kramers' rate[4], $r \propto e^{-\Delta E_b / D_\sigma}$, so holding the label over a commit window with failure probability at most $\varepsilon$ demands a minimum barrier, and — since each rejected perturbation dissipates the barrier it climbs back down — a minimum energy per invariant commit,

$$E_{\text{inv}} \;\geq\; \Delta E_b^{\min} \;=\; D_\sigma \;\ln(1/\varepsilon) + \text{const} \;\xrightarrow[D_\sigma \to k_B T]{}\; \Theta\big(k_B T \ln(1/\varepsilon)\big).$$

We read this as a Landauer-analog for reproducibility: a thermodynamic lower bound on the energy required to make an output invariant against supply noise of a given amplitude, distinct from Landauer's $k_B T \ln 2$ per erased bit[1] — Landauer prices erasure, this prices reproducibility. Spend less and the basins are too shallow to reject the noise; the reproducibility is not there to buy at any lower price. Whether this is a tight bound or can be beaten by a coarse-graining cleverer than $k_B T \ln(\text{states})$ is, in our view, the open theoretical question of the area, and we state it as a conjecture supported by a scaling argument rather than as a proved theorem.

The bound has two consequences worth drawing out. The first inverts a premise the low-power field treats as settled: reversible and adiabatic computing minimize energy 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, so the substrate that earns crossing-free invariance is minimally dissipative but deliberately *not* reversible, sitting above the Landauer floor by exactly the barrier it needs; reproducibility and reversibility pull against each other. The attractor-network construction this rests on is classical[3]; the contribution is its use as a priced, crossing-free invariance primitive. The second consequence is the honest cost: deep, well-separated basins occupy the state range at a coarser spacing, so the number of distinguishable states per device — the information capacity — falls as the barrier rises. Free invariance is paid in bits per device, and that density–margin tradeoff is the true Pareto surface of energy-native computation.

## 8. Maturity and honest limits

Two of the claims here rest on demonstrated science and one does not, and it would be misleading to present them as equal. The crossings spectrum is a classification of existing work; the identity mechanisms it points to are demonstrated for light and biology and published in halves for the ionic carrier, and to that extent the map is solid. The invariance construct is an engineering synthesis of established parts — attractor dynamics, ensemble averaging, hysteresis — and is realizable in principle, though no integrated energy-native device has been built to the specification here. The bound is the least settled element: it is a scaling argument with a clear physical mechanism, offered as a conjecture. The interactive companion is a first-order model that illustrates the signs and the floor, not a validated simulator, and the report claims no measured energies. Read together, the maturity is that of a well-grounded research position: a map that is real, a construct that is buildable, and a bound that is plausible and worth proving or refuting.

## 9. Conclusion

Energy conversion and computation converge, at the low-power limit, onto one physical mechanism, and the honest way to organize that convergence is by the crossings that survive between a machine's energy and its state. Collapsing those crossings is the efficiency win, and it is also, along the same diagonal, the reproducibility problem, because the interface you delete is the interface you would have made the computation repeatable at. The way out is not to convert and round but to let dissipation build the basins that reject supply noise for free and to commit the basin label rather than the microstate — and the price of doing so has a thermodynamic floor, a Landauer-analog for a reproducible bit. For Physical AI, which must be both frugal enough to run in a body and reproducible enough to be trusted with one, that floor is the quantity worth knowing. The map is offered as an open commons, and the bound as a question the report would rather see settled than owned.

## 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.
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. Anionic entanglement-induced giant thermopower in ionic thermoelectric materials. *eScience* 3, 100169, 2023. doi:10.1016/j.esci.2023.100169.
8. Nanofluidic logic with mechano–ionic memristive switches. *Nature Electronics* 7, 271–278, 2024. doi:10.1038/s41928-024-01137-9.
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 report used a large language model (Claude, Anthropic) for drafting and editing text, organizing the reviewed literature, and preparing the figures and 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

Technical Report TR-2026-05 · Preprint v1