Conservation-Congruent Encodings

Peter David Fagan, School of Informatics, University of Edinburgh

April 24, 2026

Abstract

A conservation-congruent encoding (CCE) is a physically realized macroscopic distinction, unlike the abstract, substrate-independent notion of information used in traditional information theory. Under a chosen coarse-graining, it is represented by protected macroscopic regions and their associated world-tubes that persist under ambient fluctuations, are maintained by dynamical invariants tied to conserved quantities, and can be irreversibly merged only through dissipative export into explicitly modeled channels. This note gives a minimal definition of a CCE.

1 Introduction

Classical information theory often treats information as a substrate-independent abstraction. In practice, however, information is always realized as a physical distinction carried by matter. Landauer's analysis of logical irreversibility and Bennett's reversible-computation perspective make clear that persistence, reversible transport, and erasure have different physical statuses [5, 1]. The relevant question is therefore how to relate information processing to descriptions of matter. We do so by considering how macroscopic distinctions can persist against fluctuations and be irreversibly merged only by paying a physically measurable cost. The Conservation-Congruent Encoding (CCE) framework developed in this note, and discussed in a companion preprint [3], addresses that question by defining information directly in terms of coarse-grained dynamics.

A physical distinction must first persist long enough to remain a stable macroscopic distinction. Once such a distinction exists, the dynamics may carry it within a protected region or merge it irreversibly across a boundary in ways consistent with the conservation structure of the substrate. The framework therefore distinguishes persistence from the reversible and irreversible processes that act on it.

That split is captured by metriplectic dynamics [6, 7]. The reversible sector transports a state within a protected region without destroying its identity, while the dissipative sector is what makes a many-to-one destruction of distinctions physically irreversible and supplies the associated physical exhaust. The rest of this note states the minimal mathematical conditions under which a macroscopic distinction qualifies as a conservation-congruent encoding.

2 Metriplectic Dynamics of Information

Let X denote the microscopic dynamical state space of the substrate, with microscopic state $x(t) \in X$ evolving under the fine-grained flow

$$\dot{x} = F(x). \tag{1}$$

Let the chosen coarse-graining define a macroscopic state space $\mathcal{Z}$ of slow, operationally resolvable variables. The projection

$$\Pi : X \to \mathcal{Z}, \qquad Z = \Pi(x), \tag{2}$$

identifies which physical distinctions count as information. When several conserved-quantity channels are operationally relevant, the same construction may be multiplexed as $\Pi = (\Pi^{(1)}, \ldots, \Pi^{(C)})$ with $\mathcal{Z} = \mathcal{Z}^{(1)} \times \cdots \times \mathcal{Z}^{(C)}$. Two microscopic states are equivalent when they map to the same macroscopic value of Z, so a CCE is defined not on X alone but on the pair (X, Π) together with the protected macroscopic regions induced by that projection. The fast degrees of freedom discarded by Π are not removed physically; they form the unresolved bath against which the retained macroscopic variables exchange conserved quantities.

For any smooth macroscopic observable $A(Z)$, the effective dynamics are written in metriplectic form,

$$\dot{A} = \{A, H\} + (A, S), \tag{3}$$

and in particular for the macroscopic coordinates,

$$\dot{Z}^a = \{Z^a, H\} + (Z^a, S), \tag{4}$$

or compactly $\dot{Z} = \{Z, H\} + (Z, S)$. Here H is the effective conserved generator and S is an entropy-like potential that orders irreversible contraction. The Poisson term describes the reduced closed dynamics on the retained macroscopic manifold, whereas the metric term collects the open-system influence of the eliminated degrees of freedom.

The reversible sector is encoded by the Poisson bracket

$$\{A, B\} = \partial_a A \, J^{ab}(Z) \, \partial_b B, \qquad J^{ab} = -J^{ba}, \tag{5}$$

with J the antisymmetric Poisson tensor. The irreversible sector is encoded by the symmetric bracket

$$(A, B) = \partial_a A \, M^{ab}(Z) \, \partial_b B, \qquad M^{ab} = M^{ba}, \quad M \succeq 0, \tag{6}$$

with M the positive semi-definite metric operator. In this form the symmetric bracket is the coarse-grained bookkeeping device for friction, relaxation, and export into the environment.

The standard metriplectic degeneracy conditions are

$$\{S, A\} = 0, \qquad (H, A) = 0, \tag{7}$$

for admissible observables A. These are independent axioms: $\{S, A\} = 0$ makes the reversible sector isentropic, while $(H, A) = 0$ ensures that the dissipative sector preserves the conserved generator. Together they imply

$$\dot{H} = 0, \qquad \dot{S} = (S, S) \geq 0. \tag{8}$$

This is the minimal dynamical split required for information: the Poisson sector preserves the invariants that stabilize an encoding, while the metric sector accounts for dissipative coarsening and exported exhaust when control protocols drive a many-to-one logical merger. Physically, the phase-space volume lost from the retained description does not disappear; it leaves the coarse-grained dynamics through the symmetric bracket and is exported into environmental degrees of freedom excluded by Π. The monotonicity $\dot{S} = (S, S) \geq 0$ is therefore a coarse-grained signature of dissipative exhaust compatible with Landauer-type lower bounds once a reservoir model, control protocol, and reset map are specified [5, 8]. In a thermal reservoir that exhaust may appear as heat, in a spin reservoir it may appear as angular momentum, and in other substrates it may appear as the corresponding conserved quantity exported against its conjugate generalized force.

When several channels are active simultaneously, the same metriplectic structure applies on the multiplexed space $\mathcal{Z} = \mathcal{Z}^{(1)} \times \cdots \times \mathcal{Z}^{(C)}$. The metric operator may be block-diagonal when channels erase independently or mixed when a single logical distinction is stabilized by coupled dynamics. In that setting the abstract geometry does not by itself fix a unique apportionment of the distinguishability loss across channels: the channel shares depend on the specified reservoir model, control protocol, and channel-specific bookkeeping factors, rather than being tied to heat alone.

3 The Geometric Conditions for a Conservation-Congruent Encoding

Under a chosen coarse-graining Π, a family of protected regions $\{\mathcal{M}_i(t)\}$ and associated world-tubes $\{\mathcal{W}_i\}$ qualifies as a Conservation-Congruent Encoding if and only if it satisfies three defining conditions. For orientation, the three conditions provide a compact roadmap from encoding to reversible manipulation to irreversible erasure:

Condition I: Metastable Confinement
Protected macroscopic cross-sections must assemble into a world-tube that survives throughout the operational window with exponentially suppressed escape.

Condition II: Transverse Restoring Structure
Each protected region must admit local charts whose transverse restoring geometry makes quasistatic deformations well posed and drives the excess channel expenditure to zero in the quasistatic limit.

Condition III: Irreversible Tube Merger
A logically irreversible operation is physically realized only when two or more disjoint world-tubes are collapsed into a common output tube, with the metric sector supplying the dissipative export that makes the merger irreversible.

Individual macroscopic states $Z \in \mathcal{M}_i(t)$ are representatives of the encoding, not the encoding by themselves. To make the temporal aspect explicit, define the extended macroscopic space

$$\mathcal{E} = \mathcal{Z} \times \mathcal{T}, \tag{9}$$

with coordinates (Z, t). In this geometric view, a logical state is not merely a basin in $\mathcal{Z}$ at an instant, but a protected world-tube

$$\mathcal{W}_i = \bigcup_{t\in[t_0,t_0+\tau_{\text{op}}]} (\mathcal{M}_i(t) \times \{t\}) \subset \mathcal{E}, \tag{10}$$

whose lateral boundary is the time-extended separatrix $\bigcup_{t\in[t_0,t_0+\tau_{\text{op}}]} (\partial\mathcal{M}_i(t) \times \{t\})$. The stationary encoding case is recovered by $\mathcal{M}_i(t) \equiv \mathcal{M}_i$. A CCE persists when the trajectory reaches the terminal slice $t_0 + \tau_{\text{op}}$ without fluctuating across that lateral boundary.

3.1 Condition I: Metastable Confinement

Condition I is the encoding criterion. A valid encoding is a protected family of macroscopic cross-sections $\mathcal{M}_i(t) \subset \mathcal{Z}$ whose union along $\mathcal{T}$ forms a metastable world-tube $\mathcal{W}_i \subset \mathcal{E}$. For a macroscopic trajectory $(Z(t), t)$, define the operational escape probability by

$$\begin{aligned} P^{(i)}_{\text{esc}}(\tau_{\text{op}}) &\equiv \Pr\left(\exists\, t \in (t_0, t_0 + \tau_{\text{op}}) : (Z(t), t) \notin \mathcal{W}_i \mid Z(t_0) \in \mathcal{M}_i(t_0)\right), \\ P^{(i)}_{\text{esc}}(\tau_{\text{op}}) &\lesssim \tau_{\text{op}}\, \nu_{0,i} \exp\left(-\frac{\Delta\mathcal{B}_i}{\sigma_i}\right), \qquad \tau_{\text{op}} \ll \tau_{\text{escape}}. \end{aligned} \tag{11}$$

Here $\nu_{0,i}$ is the characteristic intrawell attempt frequency, $\Delta\mathcal{B}_i$ is the channel-resolved quasi-potential barrier along the dominant escape path, and σ_i is the corresponding fluctuation scale. In the weak-noise regime of a stochastic reduction with Kramers-type escape one has $\tau_{\text{escape}}^{-1} \sim \nu_{0,i} \exp(-\Delta\mathcal{B}_i/\sigma_i)$, so the expression above is a model form rather than a universal consequence of coarse-graining. For nonthermal channels, Eq. (11) should therefore be read as a model-dependent weak-noise first-passage estimate: both the exponential rate function and the prefactor depend on the specified reservoir and fluctuation structure, and need not reduce to a literal energy barrier divided by k_BT. The Arrhenius–Kramers thermal limit is recovered by $\Delta\mathcal{B}_i = \Delta E_i$ and $\sigma_i = k_BT$. When the eliminated fast variables admit a near-equilibrium fluctuation-dissipation closure, M^{ab} plays the role of a mobility tensor and the fluctuation-dissipation theorem gives an Einstein relation of the form $D^{ab} = k_BT\, M^{ab}$ (up to the normalization conventions used for S), so the same mobility/diffusion data that govern Kramers escape also set the strength of metric dissipation. Outside that regime the relation is model dependent.

The tube length along the time axis is the operational window. Inside $\mathcal{W}_i$, the state may drift, circulate, or evolve on a driven invariant set without loss of logical identity. What matters is confinement until the trajectory reaches the terminal slice. Writing τ_{escape} for the mean first-passage time across the relevant tube boundary, and τ_{relax} for the local relaxation time of the fast variables eliminated by Π back toward the protected manifold, the defining timescale separation is

$$\tau_{\text{relax}} \ll \tau_{\text{op}} \ll \tau_{\text{escape}}. \tag{12}$$

3.2 Condition II: Transverse Restoring Structure

Condition II is the reversible-manipulation criterion. Condition I says that a macroscopic distinction survives for the operational window. Condition II adds that, along the operating portion of each world-tube, there is an atlas of local charts in which the protected section is transversely restoring and can be advected without loss of logical identity. Concretely, for each reference point $Z_*(t) \in \mathcal{M}_i(t)$ there must exist a chart

$$\psi_k : U_k \to \mathbb{R}^{d_t} \times \mathbb{R}^{d_n}, \qquad u = \psi_k(Z) = (q, r), \tag{13}$$

where q parameterizes directions tangent to the intended transport and r parameterizes directions normal to the tube, together with an effective quasi-potential $\Phi_k(u; \lambda)$ satisfying

$$\nabla_r \Phi_k(q, 0; \lambda) = 0, \qquad \nabla_{rr}^2 \Phi_k(q, 0; \lambda) \succeq \kappa_k I, \qquad \kappa_k > 0. \tag{14}$$

This is a local, transverse convexity requirement. Along tangent directions the tube may translate, shear, or slowly reshape under slow deformation, while in the normal directions there must remain a restoring structure that keeps nearby trajectories slaved to the protected section. Global non-convexity is therefore not excluded; indeed it is what allows distinct logical world-tubes to coexist.

Let $\lambda(t)$ be a deformation schedule that moves the reference section through this atlas without destroying the positivity of the transverse Hessian. If the schedule is written as a fixed-shape family $\lambda_{\tau_{\text{op}}}(t) = \lambda(t/\tau_{\text{op}})$, the quasistatic limit is obtained by stretching the duration rather than changing the profile itself. In that limit,

$$\sup_{t \in [0, \tau_{\text{op}}]} \left\| \dot{\lambda}_{\tau_{\text{op}}}(t) \right\| = O(\tau_{\text{op}}^{-1}), \qquad \tau_{\text{relax}} \ll \tau_{\text{op}}, \tag{15}$$

so the state tracks the moving protected section and the excess conserved-quantity expenditure above the reversible transport limit vanishes:

$$C_{\text{ex}}[\lambda_{\tau_{\text{op}}}] \to 0 \qquad \text{as} \qquad \tau_{\text{op}} \to \infty. \tag{16}$$

When linear-response assumptions apply, one expects the familiar scaling $C_{\mathrm{ex}} = O(\tau_{\mathrm{relax}}/\tau_{\mathrm{op}})$. Condition II therefore does not assert that the full logical landscape is convex. It asserts that each protected tube admits a local atlas in which reversible transport is well posed: transverse restoring geometry prevents logical leakage, and sufficiently slow deformation suppresses excess dissipation.

3.3 Condition III: Irreversible Tube Merger

Condition III is the irreversible-tube-merger criterion for erasure. Geometrically, it begins when the effective basin structure in $\mathcal{E}$ is deformed so that two or more disjoint input world-tubes can no longer remain separate protected classes and instead terminate in a common output tube. In the extended-space picture this is the analogue of a pair-of-pants manifold: distinct legs in the past join into one trunk in the future. This intended merger should be distinguished from the noise-induced escape of Condition I. During a reset, crossing the former separatrix may be part of the intended deformation because the boundary itself has been lowered, displaced, or removed.

On any constant-time slice, such a collapse identifies previously distinct macroscopic classes and therefore reduces coarse-grained distinguishability. If the input partition has Shannon entropy $\mathcal{H}_{\mathrm{in}}$ and the output partition induced by this merger has entropy $\mathcal{H}_{\mathrm{out}}$, the corresponding coarse-grained Shannon-entropy drop in nats is

$$\Delta\mathcal{H}_{\mathrm{cg}} = \mathcal{H}_{\mathrm{in}} - \mathcal{H}_{\mathrm{out}}. \tag{17}$$

For the special case of a complete merger of m input classes into a single output class with prior probabilities $\{p_i\}_{i=1}^{m}$, this reduces to

$$\Delta\mathcal{H}_{\mathrm{cg}} = -\sum_{i=1}^{m} p_i \ln p_i. \tag{18}$$

This geometric collapse by itself does not yet specify the physical expenditure. The control protocol determines how basin boundaries are removed or reconnected, while the metric sector specifies the dissipative channel through which the relevant conserved quantity is exported, thereby making the many-to-one merger physically irreversible. If a single channel c is active, let α_c denote the channel-specific bookkeeping factor that converts one nat of coarse-grained entropy drop into the corresponding channel-resolved increment, and let $\mathcal{F}_c$ be the conjugate generalized force. In the ideal symmetric case where the logical classes differ only by coarse-grained distinguishability and do not carry an additional baseline offset in the corresponding channel potential, the lower bound takes the form

$$\mathcal{J}_{c,\min}^{\mathrm{CCE}} = (\mathcal{F}_c\alpha_c)\Delta\mathcal{H}_{\mathrm{cg}}. \tag{19}$$

More general asymmetric encodings, or protocols whose endpoints differ by a nonzero baseline free-energy or generalized-work offset, require the corresponding model-dependent baseline term in addition to Eq. (19). When several channels are active simultaneously, the abstract CCE framework does not determine a unique apportionment rule by itself. A model-dependent bookkeeping parameterization is

$$\mathcal{J}_{\min}^{\mathrm{CCE}} = \sum_{c=1}^{C} \eta_c(\mathcal{F}_c\alpha_c)\Delta\mathcal{H}_{\mathrm{cg}}, \qquad \eta_c \geq 0, \qquad \sum_{c=1}^{C} \eta_c = 1, \tag{20}$$

where the channel shares η_c are fixed only after specifying the coupled reservoir model and control protocol. Without that extra structure, the multi-channel discussion and Table 1 should be read as schematic bookkeeping rather than as a predictive lower-bound formula. For an unbiased bit reset quasistatically in a single thermal channel, $\eta_1 = 1$, $\mathcal{F}_1 = T$, $\alpha_1 = k_B$, and $\Delta\mathcal{H}_{\text{cg}} = \ln 2$, so this ideal lower bound reduces to the usual Landauer bound $Q_{\min} = k_B T \ln 2$ [5]. The label "conservation-congruent" means that the encoding geometry, including the stability of world-tubes and the permitted surgery that merges them, is stabilized by and mathematically coupled to the same conservation structure that governs this export. Taken together, the three conditions are hierarchical: Condition I identifies when a distinction physically persists, Condition II identifies when it can be transported quasistatically within a transversely restoring atlas with vanishing excess cost, and Condition III identifies when that distinction is deliberately collapsed and must export the relevant conserved quantity to the environment.

4 Illustrative Physical Realizations

To anchor the definition in concrete hardware, this section pairs a canonical single-bit realization with a cross-substrate bookkeeping summary. A colloidal bead trapped in a tunable optical double-well potential and immersed in fluid provides the thermal single-channel case used in experimental Landauer tests [2, 4]. Table 1 then summarizes illustrative ways in which the same CCE geometry can be paired with different conserved quantities, generalized forces, and bookkeeping factors across superconducting, spintronic, electrochemical, and hybrid devices.

In the colloidal realization, the physical bit is carried by the bead's coarse-grained position Z in the optical landscape. The microscopic state space X is the phase space of the bead together with the surrounding water molecules, while the coarse-graining Π traces out the fluid and reduces the retained macroscopic description to the one-dimensional space $\mathcal{Z}$ coordinatized by Z. The extended macroscopic space is therefore $\mathcal{E} = \mathcal{Z} \times \mathcal{T} \cong \mathbb{R} \times \mathcal{T}$.

With the trap configured as a symmetric double-well, the left and right minima define two protected regions separated by a barrier $\Delta\mathcal{B}_i$. In the extended space these appear as disjoint world-tubes $\mathcal{W}_0$ and $\mathcal{W}_1$: thermal noise drives stochastic motion within each well, but when $\tau_{\text{op}} \ll \tau_{\text{escape}}$ the bead remains confined with exponentially suppressed escape. The local convexity of each well supplies the transverse restoring geometry required by Condition II and allows slow control deformations to transport the bead with negligible excess expenditure.

For the logically irreversible operation `RESET TO 0`, an external control $\lambda(t)$ lowers the central barrier, tilts the landscape, and then restores the double-well so that both initial logical states end in the left well. Geometrically, this executes Condition III by bending the right world-tube across the spatial axis and merging it into the left tube, yielding the familiar inverted "pair of pants" topology shown in Figure 1. Because Liouville incompressibility applies to the underlying microscopic flow, this merger is a compression only in the coarse-grained description; the surrounding fluid supplies the dissipative coupling represented phenomenologically by M^{ab}. In the idealized quasistatic isothermal limit, the mean exported heat approaches $Q_{\min} = k_B T \ln 2$ [5, 2, 4], but this thermal bound depends on the reservoir model rather than following from topology alone.

Table 1 makes the broader point explicit: the framework is universal in form but not substrate-free in content. Different devices realize the same pattern of metastable encoding, quasistatic transport, and dissipative merger through different conserved quantities, generalized forces, and bookkeeping factors, so the table should be read as an illustrative schematic rather than as a predictive lower-bound formula or a substitute for device-specific thermodynamic models. In particular, multiplexed electrochemical cases may be summarized more faithfully by a reservoir-specific work expression such as $\sum_i \Delta\mu_i \, \Delta N_i$ than by forcing every channel into a single universal one-factor template.

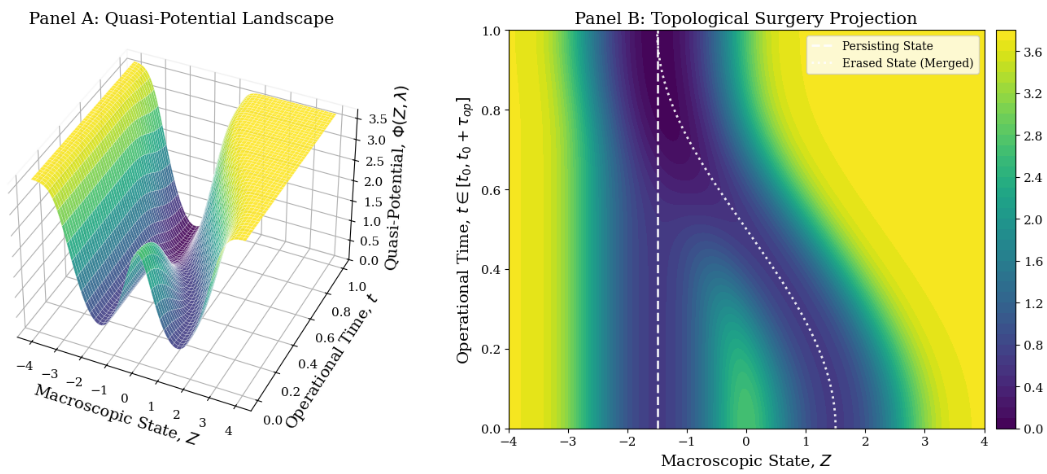

Figure 1: Pair-of-pants topology for colloidal-bit erasure in the extended macroscopic space $\mathcal{E} = \mathcal{Z} \times \mathcal{T}$, with Z the coordinate on the one-dimensional space $\mathcal{Z}$. The two disjoint input world-tubes $\mathcal{W}_0$ and $\mathcal{W}_1$ persist during encoding, then a control protocol bends the right tube across the spatial axis and merges it into the left tube, yielding a single output tube and the many-to-one geometry of `RESET TO 0`.

Table 1: Illustrative cross-substrate generalized-force and bookkeeping factors for conservation-congruent encodings, including multiplexed cases.

Encoding class	Active channel(s)	Generalized force(s) / bookkeeping factor(s)	Representative example
Thermal encoding	Energy exchanged with a thermal bath	T with entropy factor k_B	Two-state thermal encoding in a dissipative bath
Superconducting flux encoding	Magnetic flux / circulating current	I with flux increment Φ_0	Fluxoid states in a SQUID or superconducting loop
Spin or rotational encoding	Angular momentum	ω with angular-momentum increment $\hbar$	Single-domain nanomagnet or spin-torque device
Electrochemical encoding (multiplexed)	Electric charge and ionic particle-number channels	V with charge increment q; reservoir-specific chemical-work term $\sum_i \Delta\mu_i \, \Delta N_i$ (not a single-factor template)	Ion- and voltage-stabilized membrane states
Flux-spin hybrid encoding (multiplexed)	Magnetic flux and angular momentum	I with Φ_0; ω with $\hbar$	Superconducting-spin hybrid circuits or magnetic Josephson structures

5 Conclusion

This note has proposed conservation-congruent encodings as a minimal physical definition of information at the macroscopic level. In this view, information is not an abstract label floating free of matter, but a protected coarse-grained distinction specified by a substrate, a projection Π, and the metastable world-tubes that remain operationally resolvable over a finite time window. The metriplectic split provides the dynamical backdrop for that definition: the Poisson sector preserves the invariants that stabilize and transport a protected logical region, while the metric sector supplies the dissipative exhaust that accompanies irreversible many-to-one merger.

The three defining conditions summarize the framework's operational content. Metastable confinement identifies when a distinction genuinely persists, transverse restoring structure identifies when it may be manipulated quasistatically within a transversely restoring chart with vanishing excess expenditure, and irreversible tube merger identifies when distinct input world-tubes are deliberately collapsed into a common output tube so that the resulting logically irreversible map must export the relevant conserved quantity into modeled reservoir channels. Stated this way, the framework is universal in form but not substrate-free in content: different implementations stabilize and erase encodings through different conserved quantities, generalized forces, and environmental couplings.

The CCE isolates the minimal geometric and dynamical conditions under which a macroscopic distinction counts as a physically meaningful encoding, while indicating how reversible transport and irreversible merger must be grounded in conservation law.

References

[1] Charles H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–532, 1973.

[2] Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz. Experimental verification of landauer's principle linking information and thermodynamics. *Nature*, 483(7388):187–189, 2012.

[3] Peter David Fagan. Towards a physical theory of intelligence. Unpublished companion manuscript, 2026. Background companion manuscript for the broader CCE program; citation metadata provisional.

[4] Yonggun Jun, Momčilo Gavrilov, and John Bechhoefer. High-precision test of landauer's principle in a feedback trap. *Physical Review Letters*, 113(19):190601, 2014.

[5] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[6] Philip J. Morrison. Bracket formulation for irreversible classical fields. *Physics Letters A*, 100(8):423–427, 1984.

[7] Philip J. Morrison. A paradigm for joined hamiltonian and dissipative systems. *Physica D: Nonlinear Phenomena*, 18(1-3):410–419, 1986.

[8] Joanna A. Vaccaro and Stephen M. Barnett. Information erasure without an energy cost. *Proceedings of the Royal Society A*, 467(2130):1770–1778, 2011.