

Thermodynamic Efficiency Ranking of Computation-Amenable Substrates

J. R. Landers

Supplementary research note

This note supplements `dissipation-and-the-emergence-of-computation.html`, `the-geometry-of-computable-trajectories.html`, and the compact manuscript `continuous_computation.tex`. Those documents explain when a trajectory enters the computation-amenable regime and how dissipation extends the computation epoch. The present note asks a narrower comparison question: once two implementations are already in that regime, how should the repo compare their thermodynamic cheapness in a substrate-independent way?

A useful local companion is `finite-horizon-geometry-of-computational-amenability.md`, which turns the finite-horizon error budget into an observer-indexed viability region in $(\varepsilon, \eta, \lambda)$ -space. The ranking below sits on top of that geometry: it compares viable triples rather than redefining viability itself.

We work in the setting of the preceding papers. Let \mathcal{A} denote the set of trajectories admitting a coarse-graining Π with positive minimum dwell time, and let the local dissipation-lumpability relation from the main framework be

$$\lambda(P, \Pi) \approx \lambda_{\text{eq}}(\Pi) - c(\Pi)\sqrt{\sigma(P)}$$

in the perturbative regime. The associated critical dissipation scale

$$\sigma_c(\Pi) = \left(\frac{\lambda_{\text{eq}}(\Pi)}{c(\Pi)} \right)^2$$

is the point at which this local law predicts perfect lumpability ($\lambda = 0$) under coarse-graining Π .

Definition (Thermodynamic efficiency of a computing trajectory). Let $x \in \mathcal{A}$ be a trajectory governed by micro-dynamics P under coarse-graining Π . Define the *thermodynamic efficiency* of the triple (x, P, Π) as

$$\eta(x, P, \Pi) := \frac{1}{\sigma_c(x, P, \Pi)},$$

with the convention $\eta = \infty$ if $\lambda_{\text{eq}} = 0$ (the trajectory is lumpable at equilibrium). Higher η means cheaper symbolic maintenance: less entropy production is needed, in the local model, to reach the interior of \mathcal{A} .

Proposition (Efficiency ranking). The thermodynamic efficiency induces a total preorder on the set of computing triples $\{(x, P, \Pi) : x \in \mathcal{A}\}$. Explicitly:

- (i) (x_1, P_1, Π_1) is *more efficient* than (x_2, P_2, Π_2) if $\sigma_c(x_1, P_1, \Pi_1) < \sigma_c(x_2, P_2, \Pi_2)$.
- (ii) The ranking is substrate-independent: it compares across physical implementations (semiconductor, superconducting, optical, and so on) using entropy production as the common currency.
- (iii) The ranking is compatible with the dwell-time stratification of \mathcal{A} in the following sense: for fixed τ_{\min} , the most efficient trajectory is the one that achieves that dwell time at the lowest σ_c .

Definition (Constrained efficiency). For a target minimum dwell time τ^* and a class of physically realizable dynamics \mathcal{P} (for example, those governed by Maxwell's equations with a given material geometry), define the *optimal efficiency at scale τ^** :

$$\eta^*(\tau^*, \mathcal{P}) := \sup_{\substack{(x, P, \Pi) \\ P \in \mathcal{P} \\ \tau_{\min}(x, \Pi) \geq \tau^*}} \eta(x, P, \Pi) = \left[\inf_{\substack{(x, P, \Pi) \\ P \in \mathcal{P} \\ \tau_{\min}(x, \Pi) \geq \tau^*}} \sigma_c(x, P, \Pi) \right]^{-1}.$$

This defines a variational problem: among all trajectories achieving a given computational quality within a given physics, which one minimizes the thermodynamic cost of lumpability?

Proposition (Decomposition of σ_c). The critical dissipation decomposes as

$$\sigma_c = \frac{\lambda_{\text{eq}}^2}{c^2},$$

where $\lambda_{\text{eq}}(\Pi)$ measures how far the equilibrium dynamics is from lumpability (the *intrinsic lumpability gap*) and $c(\Pi)$ measures how efficiently dissipation reduces it (the *dissipation leverage*). A trajectory is efficient when

- (a) λ_{eq} is small, meaning the equilibrium dynamics is already nearly lumpable and requires little correction; or
- (b) c is large, meaning the geometry of the coarse-graining amplifies the effect of dissipation on lumpability.

The two factors are in general independent and physically distinct. Factor (a) depends on the potential landscape and its relation to the cell boundaries. Factor (b) depends on the alignment between the nonequilibrium driving and the intra-cell microstate distribution.

Definition (Efficiency-adjusted computation epoch). Combining the efficiency ranking with the thermodynamic computation epoch from the main paper, define

$$\mathcal{E}(x, P, \Pi, \dot{W}) := T_{\text{comp}}(\dot{W}) \cdot \eta(x, P, \Pi) = \frac{T_{\text{comp}}(\dot{W})}{\sigma_c(x, P, \Pi)},$$

which measures computation duration per unit of minimum thermodynamic cost. This is the natural figure of merit for comparing physical computing substrates: it rewards both long epochs and low critical dissipation.

Example (CMOS double-well). A CMOS inverter with supply voltage V_{DD} , threshold partition at $V_{DD}/2$, and effective barrier height ΔV has

- $\lambda_{\text{eq}} \sim e^{-\beta q \Delta V}$, exponentially small in the barrier height (in units of thermal voltage $k_B T/q$);
- $c \sim \sqrt{\beta/C}$, set by the gate capacitance and temperature.

Thus $\sigma_c \sim (C/\beta) e^{-2\beta q \Delta V}$, which is extremely small for typical operating conditions ($\beta q \Delta V \gg 1$). The induced symbolic trajectory sits deep inside \mathcal{A} and is highly efficient because the double-well structure provides near-perfect lumpability already at equilibrium. The dominant power consumption of a real CMOS circuit is *dynamic* (charging and discharging the capacitor during transitions), not the thermodynamic cost of maintaining lumpability. This matches the broader repo claim that maintaining a symbolic boundary and switching across that boundary are distinct thermodynamic expenses.

Example (Josephson junction). A resistively shunted Josephson junction with phase variable $\phi(t)$, governed by

$$\frac{\hbar}{2eR} \frac{d\phi}{dt} + I_c \sin \phi = I_{\text{bias}},$$

has washboard-potential wells separated by 2π in phase. Operating at millikelvin temperatures ($\beta \rightarrow \infty$), the equilibrium lumpability gap λ_{eq} becomes exponentially small, giving $\sigma_c \rightarrow 0$. The Josephson junction is then more efficient than CMOS not because its c is necessarily larger, but because the thermal noise floor is so much lower that the equilibrium dynamics is already almost perfectly lumpable.

Remark (Comparison across substrates). The ranking $\sigma_c^{\text{JJ}} \ll \sigma_c^{\text{CMOS}}$ quantifies the thermodynamic advantage of superconducting logic over semiconductor logic in the same units used by the main dissipation paper. The practical advantage is offset by the dissipation cost of cryogenic infrastructure, which the framework does not currently internalize. That external cost belongs to the admissible class \mathcal{P} , not to the intrinsic trajectory efficiency itself.

Open question (Variational solutions). For a given physical substrate (fixed \mathcal{P}), what trajectory shape minimizes σ_c at fixed τ_{min} ? In the CMOS case, the answer should recover the empirically optimal waveform: steep transitions (maximizing c by concentrating the transient in a narrow time window), deep flat plateaus (minimizing λ_{eq} by keeping microstates far from the cell boundary), and symmetric well structure (aligning the driving geometry with the coarse-graining). Whether this variational problem admits closed-form solutions for realistic nonlinear dynamics is open.