V1.0.0

CEI Invariance: Architectures as Equivalence Classes of Causal Density and Energy Fidelity

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Computational architectures are today chosen for a substrate by benchmarking, an expensive process that must be repeated for every substrate-architecture pair and gives no guidance before hardware exists. No existing framework predicts, from thermodynamics alone, which architecture classes a physical substrate can sustain. We propose the Causal-Energy Invariant (CEI): for a given substrate, the product of causal density ρ_t (dependent operations per volume-time) and energy fidelity η_e (energy per transition in units of $kT\ln 2$) is bounded by a substrate constant κ, derived from Landauer's principle, a substrate power-density limit, and causal-cone bounds on operation rates. Architectures with equal $\rho_t\eta_e$ form genuine equivalence classes, and the bound predicts that substrates favor architectures whose product approaches κ. We derive $\kappa = 3.483 \times 10^{29}\ \mathrm{m^{-3}\,s^{-1}}$ for silicon under an assumed areal power limit of $100\ \mathrm{W\,cm^{-2}}$ and a $1\ \mathrm{mm}$ active depth, show the causal-cone bound exceeds this by a factor $\approx 8.6 \times 10^{10}$ (so power, not relativity, binds silicon), and classify two assumed architecture profiles: a pipelined CMOS-like profile saturating 3.0% of κ and a gate-based error-corrected profile with $\eta_e \geq 100$ from published thermodynamic floors. Confirmation would replace per-pair benchmarking with a thermodynamic substrate-architecture map; falsification would localize where the tradeoff breaks.

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.). CEI Invariance: Architectures as Equivalence Classes of Causal Density and Energy Fidelity. QNFO Preprint v1.0.0.

1. Introduction

Cloud providers, defense laboratories, and embedded-systems designers face a recurring and costly decision: which computational architecture to build for which physical substrate. The workloads they care most about — combinatorial optimization, privacy-preserving computation, large-scale inference — are dominated by hardware whose cost scales with the number of discrete operations executed, and special-purpose accelerators, neuromorphic chips, and optical processors are today evaluated by qubit or transistor counts rather than by any physical measure of fit between machine and medium [13]. Every candidate architecture must be designed, fabricated, and benchmarked against every candidate substrate before a selection can be made. A failed pairing costs a fabrication cycle and years of engineering effort; the decision would change qualitatively if a substrate could be characterized once, by a small number of physical constants, and each architecture classified against it on paper before any hardware is built.

The gap this paper addresses is the absence of a thermodynamic quantity that (i) is a property of the substrate alone, (ii) is a property of the architecture alone, and (iii) constrains their combination. Energy-per-operation measures and operation-rate measures exist separately, and Landauer's principle gives a floor for the former [13], but no published framework composes them into a single invariant that partitions architectures into equivalence classes per substrate. This paper proposes such a quantity, the Causal-Energy Invariant (CEI), defined as the product

$$\Pi \;=\; \rho_t\,\eta_e,$$

where ρ_t is the causal density — the number of mutually dependent operations executed per unit volume per unit time — and η_e is the energy fidelity — the energy dissipated per elementary transition divided by the Landauer limit $kT \ln 2$. The central claim is that for a substrate with a finite power-density budget, $\rho_t \eta_e$ cannot exceed a substrate constant κ, and that architectures saturating or approaching κ on the same substrate form an equivalence class of thermodynamically equivalent designs. The claim is deliberately stated so that it can be falsified: if measured platforms at their power limits show products differing by more than an order of magnitude on the same substrate, the invariant fails, and the failure localizes which assumption (Landauer floor, power budget, or causal cone) is responsible.

The claim matters beyond computer engineering. Equivalence-class methods have organized models in knot theory, statistics, graph polynomials, and signal processing; a thermodynamic equivalence relation over architectures would give physical computing a comparable classification tool, and the same product form could be tested on biological and engineered systems whose equivalence has been argued on categorical grounds [12].

The contributions of this paper are:

- A derivation of the CEI bound κ from Landauer's principle, a substrate power-density limit, and causal-cone bounds on operation rates, with all arithmetic shown (Section 4).
- A proof that the relation "equal $\rho_t \eta_e$" is a genuine equivalence relation and that its classes are the orbits of a one-parameter partition, following the methodological standard for equivalence-class constructions (Section 4).
- A worked numerical classification for silicon under explicitly stated assumptions, giving $\kappa = 3.483 \times 10^{29}\ \mathrm{m^{-3}\,s^{-1}}$, a causal-to-power bound ratio of 8.607×10^{10}, and two architecture placements at 3.0% of κ and, in principle, at the bound itself ($\Pi_G / \kappa \leq 1$) (Section 5).
- A mismatch-penalty projection for running an architecture on a substrate whose κ is smaller than the architecture's product, and explicit falsification criteria with the observations that would refute the main result (Sections 5 and 6).

The product invariant, if confirmed on measured platforms, would also bear on the classification of biological computation, where energy management and information processing are already argued to be isomorphic across domains [12], and on network-level comparisons of computing systems, where equivalence is currently assessed across separate structural and dynamic layers rather than by a single composed quantity [14].

2. Background and Related Work

Equivalence classes are a mature organizing device across mathematics and its applications, and the present paper imports their methodological standards into the thermodynamics of computing. This section reviews the specific results of the cited works that the construction relies on.

Nahas [1] studies equivalence classes of Fox m-colorings of link diagrams, where a coloring assigns integers modulo m to arcs subject to the crossing relation $2b - a$, and defines classes as orbits of coloring automorphisms — permutations f satisfying $f(a * b) = f(a) * f(b)$ for the dihedral quandle operation $x * y = 2y - x \pmod m$. Its Proposition 3.1 gives a closed count: a link with p-nullity n for odd prime p has $(p^{n-1} - 1)/(p - 1)$ equivalence classes of p-colorings. The lesson imported here is that an equivalence class construction is complete only when the classes can be counted or parameterized by an invariant; Section 4 parameterizes CEI classes by the single scalar Π in the same spirit.

The chain event graph literature [2] characterizes statistical equivalence classes of staged trees by a shared polynomial description: two staged trees are statistically equivalent when their parametrization rules — products of transition probabilities along root-to-leaf paths — parametrize the same model. That work notes that a polynomial equivalence class contains $2^{\#\text{twins}}$ elements from naive swaps, and that a naïve count of the elements in the class it analyzes reveals

nearly one thousand elements, of which only a fraction correspond to staged trees. This warns that a class defined by an invariant may contain many members that are not admissible instances of the structure; Section 6 carries the same caution over to architectures that share a Π value but are not physically realizable on the target substrate.

Aiken and Ng on independence equivalence of cycles [3] and of paths [6], and Brown and Alikhani on domination equivalence of paths [10], solve the analogous problem for graph polynomials: two graphs are equivalent when they share the same independence or domination polynomial, and the works determine exactly which non-isomorphic graphs share a polynomial. [3] records that the roots of $I(C_n, x)$ are $c_i = -1/(2 + 2\cos((2i-1)\pi/n))$, and [6] gives the corresponding roots $p_i = -1/(2 + 2\cos(2i\pi/(n+2)))$ for paths, showing that equivalence classes can be read off a factorization of the defining invariant. The structural moral — that the invariant's factorization determines the class membership — motivates Section 4's factorization of Π into a substrate factor and an architecture factor.

The formalization of quotients in [4] treats the computational content of equivalence classes: defining a function on a quotient requires proving the result independent of the chosen representative, a proof obligation that users of theorem provers discharge by preferring canonical representatives such as de Bruijn variables for α-equivalence. Section 4 discharges the analogous obligation for CEI by showing that the mismatch penalty of Section 4.6 depends only on the product Π, not on the separate factors.

Pigolotti, Roldán and collaborators on large deviations [5] define an equivalence relation on stochastic dynamics under which dynamics in the same class share the same large-deviation properties; dynamics factor into an equilibrium part (common to the class) and a nonequilibrium part carrying the affinities, decomposed through the operators $E_{ij}[P] = \sqrt{P_{ij}P_{ji}}$ and a chord-based $Z(\boldsymbol{\mu})$. This is the closest structural precedent to CEI: a thermodynamic quantity (the large-deviation rate function) is constant across an equivalence class of physical processes. CEI applies the same pattern with the product $\rho_t \eta_e$ in place of the rate function.

Calimani and collaborators on Cayley graphs [9] characterize weak equivalence of Cayley graphs $C(\mathcal{A}, \Omega)$ by the existence of a group automorphism α with $\alpha(\Omega) = \Omega'$, and count weak equivalence classes as orbits of the $\mathrm{Aut}(\mathcal{A})$ action on generator sets. The orbit-counting technique — classes as orbits of a symmetry group acting on the objects — is the formal template Section 4 uses when it observes that rescaling (ρ_t, η_e) by $(s, 1/s)$ leaves Π fixed, so classes are orbits of a one-parameter rescaling action.

Segarra and Marques on graph Fourier transforms [8] define isomorphic and Jordan equivalence classes over which a spectral projector-based graph Fourier transform is identical, showing that the transform can be computed once per class rather than once per graph, with Jordan classes permitting identical transforms over graphs of nonidentical topology. This is the practical payoff pattern CEI aims at: computation done once per equivalence class instead of once per instance — here, substrate characterization done once instead of benchmarking per architecture-substrate pair.

Among works directly on physical computation, the QNFO categorical-equivalence paper [12] argues that biological and engineered computational systems implement categorically equivalent architectures through mathematical isomorphisms in energy management, information processing, and topological protection, and proposes multi-tiered experimental validation including two-dimensional electronic spectroscopy and neuromorphic benchmarking against conventional architectures. CEI supplies a candidate quantitative invariant that such a validation program could measure; the two programs are complementary, with [12] asserting the equivalence exists and CEI specifying a scalar to test it on.

The Hamiltonian-isomorphism taxonomy [13] is the nearest quantitative predecessor. It defines Hamiltonian isomorphism as an equivalence relation between a problem's cost structure and a substrate's energy function, proves it reflexive, symmetric, and transitive, and derives thermodynamic floors with all arithmetic shown: a single-bit erasure floor of 2.871×10^{-21} J at 300 K from $kT \ln 2$; a gate-based logical-operation floor of 2.871×10^{-19} to 2.871×10^{-18} J for error-correction overheads $N_{\text{ec}} \in [10^2, 10^3]$; a single-relaxation machine floor of 6.029×10^{-20} J for $n_{\text{out}} = 20$; and an energy ratio of 4.762×10^6 for a 10^6-gate instance against a 10^2-overhead gate machine. Its discussion also

catalogs the costs its floor assumption omits: analog implementation dissipation, Hamiltonian programming cost, and readout noise forcing repetition. CEI adopts its Landauer arithmetic and its floor-not-devices discipline, and extends the classification from problem-substrate pairs to architecture-substrate classes.

Finally, the network-isomorphism framework [14] argues that equivalence between complex systems cannot be captured by a single measure but must be assessed progressively across layers — topological structure, functional architecture, and dynamic behavior — formalized as an L1-L2-L3 hierarchy with an Abstract Network Object as the comparison tool. CEI is deliberately a single-layer, single-scalar proposal; Section 6 discusses what the L1-L2-L3 lesson implies, namely that a one-scalar invariant will miss equivalences and inequivalences that only appear at the dynamic layer.

3. Methods

This paper is analytical: it defines a model of substrate-architecture interaction, derives the CEI bound from three physical inputs, proves the associated relation is an equivalence relation, and computes numerical classifications under explicitly stated assumptions. No experiment, measurement, or simulation on hardware was performed; every number is either derived in Section 4 with shown arithmetic, adopted from a cited work with attribution, or labeled an assumption.

Model. A computational substrate is a physical medium with (i) an areal power-dissipation limit $p_{\max}$ in $\mathrm{W\,m^{-2}}$, (ii) an active depth ℓ in m over which dissipation occurs, (iii) a minimum feature length L in m, and (iv) an operating temperature T. An architecture on that substrate is characterized by two numbers: the causal density

$$\rho_t \;=\; \frac{\text{number of mutually dependent operations}}{\text{volume} \times \text{time}} \quad [\mathrm{m^{-3}\,s^{-1}}],$$

and the energy fidelity

$$\eta_e \;=\; \frac{E_{\mathrm{tr}}}{kT\ln 2} \quad [\text{dimensionless}],$$

where E_{tr} is the mean energy dissipated per elementary transition. Dependent operations are those that cannot begin before their predecessors complete, so ρ_t counts only causally coupled computation; independent parallel lanes do not increase ρ_t beyond the dependency-limited rate.

Assumptions. We assume: (A1) every elementary transition dissipates at least the Landauer limit $kT\ln 2$, the standard thermodynamic floor for irreversible state change, consistent with the 2.871×10^{-21} J floor at 300 K derived in [13]; (A2) the substrate cannot sustainably dissipate more than $p_{\max}$ per unit area, so the volumetric budget is $p_{\max}/\ell$; (A3) no causal density can exceed the relativistic bound set by the feature length, $\rho_t \le c/L^4$, because a dependent operation needs at least one signal traversal of length L (time $\ge L/c$) within a volume of at least L^3; (A4) for the worked silicon example, $p_{\max} = 100\ \mathrm{W\,cm^{-2}}$, $\ell = 1\ \mathrm{mm}$, $L = 10\ \mathrm{nm}$, $T = 300\ \mathrm{K}$ — these are stated assumptions standing in for measured platform data, chosen as round representative values, not measurements.

Derivation and computation plan. Section 4 proceeds in four steps: (1) derive the power-budget bound on $\rho_t\eta_e$; (2) derive the causal bound and compare; (3) prove the equal-product relation is an equivalence relation and identify its classes as orbits of a rescaling action; (4) place two assumed architecture profiles on the silicon scale and derive the mismatch penalty for a throttled architecture. Section 5 reports only these computed values and clearly labeled projections.

4. Analysis

4.1 The Landauer constant

Input: Boltzmann constant $k = 1.380649 \times 10^{-23}\ \mathrm{J\,K^{-1}}$; $T = 300\ \mathrm{K}$ (Assumption A4); $\ln 2 = 0.6931472$.

$$kT \;=\; (1.380649 \times 10^{-23}) \times 300 \;=\; 4.141947 \times 10^{-21}\ \mathrm{J},$$

$$kT \ln 2 \;=\; (4.141947 \times 10^{-21}) \times 0.6931472 \;=\; 2.871 \times 10^{-21}\ \mathrm{J}.$$

This reproduces the single-bit erasure floor derived in [13], which states the same value 2.871×10^{-21} J at 300 K, computed from $kT \ln 2$. We denote $\varepsilon_L = 2.871 \times 10^{-21}$ J.

4.2 The power-budget bound on $\rho_t \eta_e$

By Assumption A1, each transition dissipates $E_{\mathrm{tr}} = \eta_e\, \varepsilon_L$. The volumetric rate of transitions is ρ_t, so the volumetric dissipated power is

$$P_V \;=\; \rho_t\, E_{\mathrm{tr}} \;=\; \rho_t\, \eta_e\, \varepsilon_L.$$

By Assumption A2, sustainability requires $P_V \le p_{\max}/\ell$. Therefore

$$\rho_t\, \eta_e \;\le\; \frac{p_{\max}}{\ell\, \varepsilon_L} \;\equiv\; \kappa.$$

The product $\Pi = \rho_t \eta_e$ is bounded by a constant κ that depends only on substrate quantities ($p_{\max}$, ℓ) and the operating temperature (through ε_L). This is the CEI bound. Its units check: $[\mathrm{W\,m^{-2}}]/([\mathrm{m}][\mathrm{J}]) = \mathrm{m^{-3}\,s^{-1}}$, matching the units of ρ_t with η_e dimensionless.

4.3 The causal-cone bound and which constraint binds

By Assumption A3, a dependent operation occupies at least a volume L^3 and at least a time L/c, so

$$\rho_t^{\mathrm{causal}} \;=\; \frac{c}{L^4}.$$

With $c = 2.998 \times 10^{8}\ \mathrm{m\,s^{-1}}$ and $L = 10\ \mathrm{nm} = 1.0 \times 10^{-8}\ \mathrm{m}$ (Assumption A4):

$$L^4 = (1.0 \times 10^{-8})^4 = 1.0 \times 10^{-32}\ \mathrm{m^4},$$

$$\rho_t^{\mathrm{causal}} = \frac{2.998 \times 10^{8}}{1.0 \times 10^{-32}} = 2.998 \times 10^{40}\ \mathrm{m^{-3}\,s^{-1}}.$$

The ratio of the causal bound to the power bound:

$$\frac{\rho_t^{\mathrm{causal}}}{\kappa} \;=\; \frac{2.998 \times 10^{40}}{3.483 \times 10^{29}} \;=\; 8.607 \times 10^{10},$$

where κ is computed in Section 4.5 below. The causal bound exceeds the power bound by a factor of approximately 8.6×10^{10}: for silicon at the assumed parameters, the power budget, not special relativity, is the binding constraint by roughly eleven orders of magnitude. The CEI bound is therefore the operative invariant on this substrate, and the causal cone enters only as a far upper ceiling.

4.4 Equivalence classes

Define on the set of architecture profiles (ρ_t, η_e) the relation

$$(\rho_t, \eta_e) \sim (\rho_t', \eta_e') \quad \Longleftrightarrow \quad \rho_t\,\eta_e = \rho_t'\,\eta_e'.$$

Reflexivity: $\rho_t\eta_e = \rho_t\eta_e$. Symmetry: equality of reals is symmetric. Transitivity: $\rho_t\eta_e = \rho_t'\eta_e'$ and $\rho_t'\eta_e' = \rho_t''\eta_e''$ imply $\rho_t\eta_e = \rho_t''\eta_e''$. Hence $\sim$ is an equivalence relation, and the classes are genuine equivalence classes in the sense required by the methodological standard that a foundational relation be a true equivalence relation, the same standard the Hamiltonian-isomorphism taxonomy verifies for its own relation [13].

The classes admit an orbit description. For $s > 0$ define the rescaling action $s \cdot (\rho_t, \eta_e) = (s\,\rho_t, \eta_e/s)$. Then $\rho_t\eta_e$ is invariant under the action, and conversely any two profiles with the same product are related by $s = \rho_t'/\rho_t$. The equivalence classes are exactly the orbits of this one-parameter action, parameterized by the scalar $\Pi \in (0, \kappa]$. This mirrors the orbit-counting structure of Cayley-graph equivalence, where classes are orbits of the automorphism group acting on generator sets [9], and the coloring-class construction, where classes are orbits of coloring automorphisms [1].

The quotient obligation of [4] — that a function on classes be independent of the representative — is discharged for the mismatch penalty of Section 4.6: the penalty depends only on the ratio Π/κ, which is constant on each class, so any representative of a class receives the same penalty.

4.5 Numerical value of κ for silicon

Inputs (Assumption A4): $p_{\max} = 100\ \mathrm{W\,cm^{-2}} = 100/(10^{-4})\ \mathrm{W\,m^{-2}} = 1.000 \times 10^{6}\ \mathrm{W\,m^{-2}}$; $\ell = 1\ \mathrm{mm} = 1.0 \times 10^{-3}\ \mathrm{m}$; $\varepsilon_L = 2.871 \times 10^{-21}\ \mathrm{J}$ from Section 4.1.

Volumetric power budget:

$$\frac{p_{\max}}{\ell} = \frac{1.000 \times 10^{6}}{1.0 \times 10^{-3}} = 1.000 \times 10^{9}\ \mathrm{W\,m^{-3}}.$$

CEI constant:

$$\kappa \;=\; \frac{1.000 \times 10^{9}}{2.871 \times 10^{-21}} \;=\; 3.483 \times 10^{29}\ \mathrm{m^{-3}\,s^{-1}}.$$

4.6 Placement of two architecture profiles

Profile S (pipelined CMOS-like). Assumptions: device density $n_d = 1.0 \times 10^{15}\ \mathrm{m^{-3}}$, clock frequency $f = 3.0 \times 10^{9}\ \mathrm{s^{-1}}$, activity factor $\alpha = 0.10$, mean energy per transition $E_{\mathrm{tr}} = 1.0 \times 10^{-16}\ \mathrm{J}$. All four are stated assumptions, not measurements.

$$\rho_t^{(S)} = n_d\, f\, \alpha = (1.0 \times 10^{15}) \times (3.0 \times 10^{9}) \times 0.10 = 3.0 \times 10^{23}\ \mathrm{m^{-3}\,s^{-1}},$$

$$\eta_e^{(S)} = \frac{1.0 \times 10^{-16}}{2.871 \times 10^{-21}} = 3.483 \times 10^{4},$$

$$\Pi_S = (3.0 \times 10^{23}) \times (3.483 \times 10^{4}) = 1.045 \times 10^{28}\ \mathrm{m^{-3}\,s^{-1}}.$$

Fraction of the bound:

$$\frac{\Pi_S}{\kappa} = \frac{1.045 \times 10^{28}}{3.483 \times 10^{29}} = 0.0300.$$

Profile S saturates 3.0% of the silicon CEI bound. Consistency check against the power budget: $\rho_t^{(S)} E_{\mathrm{tr}} = (3.0 \times 10^{23}) \times (1.0 \times 10^{-16}) = 3.0 \times 10^{7}\ \mathrm{W\,m^{-3}}$, which is 3.0% of the volumetric budget $1.000 \times 10^{9}\ \mathrm{W\,m^{-3}}$, confirming internal consistency.

Profile G (gate-based, error-corrected). [13] derives a gate-based logical-operation floor of 2.871×10^{-19} J for error-correction overhead $N_{\text{ec}} = 10^2$ and 2.871×10^{-18} J for $N_{\text{ec}} = 10^3$. Taking the lower value as the assumed transition energy:

$$\eta_e^{(G)} = \frac{2.871 \times 10^{-19}}{2.871 \times 10^{-21}} = 100.$$

The causal density of profile G is bounded by the same power budget: at $\eta_e = 100$, the maximum sustainable causal density is

$$\rho_t^{(G)} \leq \frac{\kappa}{\eta_e^{(G)}} = \frac{3.483 \times 10^{29}}{100} = 3.483 \times 10^{27}\ \mathrm{m^{-3}\,s^{-1}}.$$

Profile G therefore occupies the class $\Pi = 3.483 \times 10^{29}\ \mathrm{m^{-3}\,s^{-1}}$ at most — that is, an error-corrected gate-based architecture at the Landauer-derived floor can in principle saturate the CEI bound exactly, because its energy fidelity is fixed by the floor and its causal density adjusts to the power budget. Its placement relative to κ is thus $\Pi_G/\kappa \leq 1$, with equality achievable only if the substrate can actually be operated at its full areal power limit with no peripheral overheads.

Comparison of the two profiles. The ratio of products is

$$\frac{\Pi_G}{\Pi_S} = \frac{3.483 \times 10^{29}}{1.045 \times 10^{28}} \approx 33.3,$$

so the two assumed profiles differ by roughly one and a half orders of magnitude in their approach to the same substrate bound. Under the CEI hypothesis, they belong to distinct equivalence classes on silicon, with profile G closer to the thermodynamic ceiling.

4.7 Mismatch penalty for a throttled architecture

Consider an architecture whose natural product Π_{arch} exceeds the substrate constant κ_{sub} of the substrate it is implemented on. Sustainability (Assumption A2) forces the operator to throttle either the causal density or the energy fidelity — in practice the causal density, since η_e is fixed by the device technology. The throttle factor is

$$\theta \;=\; \frac{\kappa_{\text{sub}}}{\Pi_{\text{arch}}} \;\leq\; 1,$$

and the achievable causal density becomes $\theta\,\rho_t$, a fractional loss $(1-\theta)$ in throughput. Because θ depends only on the ratio $\Pi_{\text{arch}}/\kappa_{\text{sub}}$, the penalty is constant on each equivalence class, discharging the quotient obligation of Section 4.4.

As a worked projection: first, a substrate with an assumed areal power limit ten times lower, $p_{\max} = 10\ \mathrm{W\,cm^{-2}} = 1.0 \times 10^5\ \mathrm{W\,m^{-2}}$, and the same $\ell = 1$ mm, gives

$$\kappa' = \frac{1.0 \times 10^{8}}{2.871 \times 10^{-21}} = 3.483 \times 10^{28}\ \mathrm{m^{-3}\,s^{-1}},$$

$$\theta = \frac{3.483 \times 10^{28}}{1.045 \times 10^{28}} = 3.333.$$

Since $\theta = 3.333 > 1$, profile S still fits without throttling on this substrate: $\Pi_S = 1.045 \times 10^{28} < \kappa' = 3.483 \times 10^{28}$. To force throttling, the areal power limit must fall below $p_{\max} = 1\ \mathrm{W\,cm^{-2}} = 1.0 \times 10^4\ \mathrm{W\,m^{-2}}$, for which

$$\kappa'' = \frac{1.0 \times 10^{7}}{2.871 \times 10^{-21}} = 3.483 \times 10^{27}\ \mathrm{m^{-3}\,s^{-1}},$$

$$\theta = \frac{3.483 \times 10^{27}}{1.045 \times 10^{28}} = 0.333.$$

The throttled architecture would then lose $1 - 0.333 = 0.667$, i.e. two thirds of its causal density, purely from the power budget. This projection assumes the same active depth and temperature; it illustrates how the CEI converts a substrate change into a concrete throughput penalty.

5. Results

All numbers in this section are computed in Section 4 from the stated inputs, or are labeled projections. No measurement or simulation was performed.

Substrate constant for silicon (Assumption A4). Under $p_{\max} = 100\ \mathrm{W\,cm^{-2}}$, $\ell = 1\ \mathrm{mm}$, $T = 300\ \mathrm{K}$:

$$\kappa = 3.483 \times 10^{29}\ \mathrm{m^{-3}\,s^{-1}}.$$

Binding constraint. The causal-cone bound $\rho_t^{\text{causal}} = c/L^4 = 2.998 \times 10^{40}\ \mathrm{m^{-3}\,s^{-1}}$ exceeds κ by a factor of 8.607×10^{10}. Power, not relativity, is the binding constraint on silicon at the assumed parameters.

Equivalence structure. The relation $\rho_t \eta_e = \rho_t' \eta_e'$ is an equivalence relation; its classes are the orbits of the rescaling action $s \cdot (\rho_t, \eta_e) = (s\rho_t, \eta_e/s)$, parameterized by $\Pi \in (0, \kappa]$.

Architecture placements (assumed profiles).

Profile	ρ_t ($\mathrm{m^{-3}\,s^{-1}}$)	η_e	Π ($\mathrm{m^{-3}\,s^{-1}}$)	Π/κ
S (pipelined CMOS-like)	3.0×10^{23}	3.483×10^{4}	1.045×10^{28}	0.0300
G (gate-based, error-corrected)	$\leq 3.483 \times 10^{27}$	100	$\leq 3.483 \times 10^{29}$	≤ 1

Profile S saturates 3.0% of the silicon bound; profile G, at the gate-based floor of [13], can in principle reach the bound. The two profiles occupy distinct equivalence classes separated by a factor of approximately 33.3 in product.

Mismatch penalty (projection). Implementing profile S on a substrate with $p_{\max} = 10\ \mathrm{W\,cm^{-2}}$ (all other parameters unchanged) would force no throttling, since $\theta = 3.333 > 1$; the throttle factor $\theta = 0.333$, a two-thirds loss in causal density, arises only at $p_{\max} = 1\ \mathrm{W\,cm^{-2}}$. This projection assumes the same active depth $\ell = 1\ \mathrm{mm}$ and temperature 300 K.

Sensitivity of κ to the assumptions. κ scales linearly in $p_{\max}$, linearly in $1/\ell$, and linearly in $1/\varepsilon_L \propto 1/T$. A factor-of-two error in any of $p_{\max}$, ℓ, or T shifts κ by a factor of two. The qualitative conclusion that power binds rather than causality survives any assumption change smaller than ten orders of magnitude.

6. Discussion

What the invariant does and does not claim. The CEI is a bound, not an equality: it states that $\rho_t \eta_e \leq \kappa$, and that designs approaching κ on the same substrate are thermodynamically equivalent in the precise sense of sharing a class. It does not claim that all architectures on a substrate converge to κ, nor that a high Π is desirable independent of workload — a profile can saturate the bound while solving the wrong problem.

The single-scalar limitation. The network-isomorphism framework [14] argues that equivalence between complex systems must be assessed across topological, functional, and dynamic layers, because no single measure captures it. CEI is deliberately single-scalar, and this is its principal limitation: two architectures in the same Π class may differ arbitrarily in latency structure, fault behavior, and programmability. The invariant classifies thermodynamic fit, not computational adequacy.

Class membership versus realizability. The staged-tree equivalence work [2] shows that a polynomial equivalence class can contain many representations that are not admissible instances of the structure: in the class it analyzes, a naïve count reveals nearly one thousand elements, generated from $2^{\#\text{twins}}$ swap-derived elements together with resize operations, of which only a fraction correspond to staged trees, a check that work leaves for future implementation. The analogous failure mode for CEI is an architecture profile that shares a Π value with a good design but cannot be realized on the substrate — for example, because its device density or clock rate exceeds what fabrication allows. The CEI bound is a necessary, not a sufficient, condition for admissibility.

What the floor assumption omits. The Hamiltonian-isomorphism taxonomy [13] catalogs the costs its Landauer-floor assumption omits: analog implementation dissipation, Hamiltonian programming cost, and readout noise forcing repetition. CEI inherits these omissions. A relaxation-based or analog architecture may dissipate substantially more per transition than its thermodynamic floor, which would place its true η_e far above the floor-derived value and move it to a different equivalence class. All placements in Section 5 that use floor values are therefore upper bounds on how close an architecture can come to κ, not predictions that it will.

Falsification. The main result — that a substrate's admissible architecture products are capped by κ and partition into equivalence classes — would be falsified by a repeatable measurement of a platform operating sustainably at a product $\rho_t \eta_e$ exceeding its independently characterized $\kappa = p_{\max}/(\ell\, \varepsilon_L)$ by more than an order of magnitude. Such an observation would localize the failure: if the excess coincides with reversible or error-corrected operation, Assumption A1 (the Landauer floor per transition) is implicated; if it coincides with concentrated heat-spreading structures, Assumption A2 (the uniform areal power limit) is implicated. A second falsifier is distributional: if measured platforms on the same substrate, each at its own power limit, show saturated products differing by more than an order of magnitude, the claim that κ organizes the architectures into a narrow band fails.

Limitations list. The results rest on the following assumptions, with their sensitivity:

The results rest on the following assumptions, with their sensitivity:

- **A1 (Landauer floor per transition).** Every elementary transition dissipates at least $kT \ln 2$. If reversible or error-corrected operation can dissipate less per logical transition than the floor assigns per elementary transition, η_e is overestimated and κ underestimated; the sensitivity is linear, so a factor-α error in the floor shifts κ by $1/\alpha$.
- **A2 (uniform areal power limit).** The substrate dissipates no more than $p_{\max}$ per unit area, spread over depth ℓ. Real substrates concentrate heat; if effective ℓ is smaller than assumed, κ falls proportionally. A factor-of-two error in ℓ shifts κ by a factor of two.
- **A3 (causal-cone ceiling).** The bound $\rho_t \le c/L^4$ assumes one signal traversal per dependent operation in a volume L^3. It is a loose ceiling — the computed ratio 8.607×10^{10} shows it is not binding — so errors here do not affect the operative bound unless they exceed ten orders of magnitude.
- **A4 (silicon parameter values).** $p_{\max} = 100\ \mathrm{W\,cm^{-2}}$, $\ell = 1\ \mathrm{mm}$, $L = 10\ \mathrm{nm}$, $T = 300\ \mathrm{K}$ are round representative assumptions, not measurements of any platform. All headline numbers scale linearly with these; the qualitative conclusion (power binds, profiles occupy distinct classes) survives any parameter change smaller than ten orders of magnitude.
- **Profile parameters.** The device density, clock, activity factor, and transition energy of profile S are assumptions standing in for measured platform data; the placement at 3.0% of κ is a projection conditional on them, not a measurement of any commercial processor.
- **Single-scalar classification.** The equivalence relation captures only the product $\rho_t \eta_e$; architectures in the same class may differ in latency, fault behavior, and programmability, as the layered-equivalence argument of [14] anticipates.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The prospective users are substrate technology teams at semiconductor manufacturers, accelerator design groups, and research programs evaluating unconventional substrates (superconducting, photonic, molecular). If the CEI bound holds, the decision that changes is the substrate-characterization step: instead of benchmarking each candidate architecture on each substrate, a team would measure three substrate quantities — areal power limit $p_{\max}$, active depth ℓ, and operating temperature T — compute $\kappa = p_{\max}/(\ell\, kT \ln 2)$, and classify each candidate architecture on paper by its product $\rho_t \eta_e$. Architectures whose product exceeds κ would be rejected before fabrication, with the throttle factor $\theta = \kappa_{\text{sub}}/\Pi_{\text{arch}}$ giving a quantitative throughput penalty for forcing the pairing. The worked projection of Section 4.7 illustrates the decision content: moving profile S to a substrate with a hundredfold lower areal power limit (to $1\ \mathrm{W\,cm^{-2}}$) costs two thirds of its causal density under the stated assumptions.

(b) Current readiness. This work is purely analytical and has not been peer reviewed: under the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and reported; concept formulated). No hardware was characterized, no prototype exists, and the framework has not been applied to any measured platform. Analytical work of this kind cannot exceed TRL 3, and reaching even TRL 3 would require the proof-of-concept measurement described next.

(c) The smallest next experiment. The minimal de-risking test is a single-substrate calibration: on one CMOS test chip, measure the areal power dissipation $p_{\max}$ under sustained load, the active dissipation depth ℓ (e.g., by thermal profiling through the die stack), and the sustained dependent-operation rate and per-transition energy of one workload, then compute ρ_t, η_e, and Π and check whether $\Pi \le \kappa$ with the arithmetic of Section 4. The outcome that refutes the framework is a sustainably operated platform whose product exceeds its independently characterized κ by more than an order of magnitude; the outcome that supports it is a measured product within the bound and consistent with the power-budget identity $P_V = \rho_t \eta_e \varepsilon_L$.

(d) Binding constraints. The constraints that would bind first in any empirical program are, in the order the derivation fixes them: (i) thermal measurement fidelity — κ is linear in $p_{\max}$ and $1/\ell$, so a factor-of-two error in either propagates directly into the classification; (ii) the definition and counting of "mutually dependent operations," which determines ρ_t and is a measurement-protocol question the paper leaves open; (iii) peripheral overheads (clocking, I/O, cooling) that the Landauer-floor discipline of [13] explicitly omits and that would raise measured η_e above any floor-derived value. Whether activity-factor-style counting of dependent operations can be made reproducible across laboratories is an open question this paper does not resolve.

(e) Order-of-magnitude bounds (projections). Under the stated assumptions only: a substrate with $p_{\max} = 100\ \mathrm{W\,cm^{-2}}$ and $\ell = 1\ \mathrm{mm}$ at $300\ \mathrm{K}$ admits at most $\kappa = 3.483 \times 10^{29}\ \mathrm{m^{-3}\,s^{-1}}$ in product; an error-corrected gate-based architecture at the $2.871 \times 10^{-19}\ \mathrm{J}$ floor of [13] ($\eta_e = 100$) could in principle sustain $\rho_t \le 3.483 \times 10^{27}\ \mathrm{m^{-3}\,s^{-1}}$; a CMOS-like profile at the assumed parameters reaches 3.0% of the bound. These are projections conditional on Assumption A4 and the profile parameters, not performance figures for any existing product.

(f) Risks and misuse paths. The invariant could be misused as a marketing or procurement gate: a Π/κ fraction near unity does not certify computational adequacy, latency, or correctness, and Section 6's realizability caveat means a high product on paper may be unrealizable on the target substrate. Decisions based on CEI alone, without the layered checks [14] recommends, would risk selecting thermodynamically efficient but functionally unsuitable architectures.

8. Conclusion

This paper proposed the Causal-Energy Invariant: for a substrate with areal power limit $p_{\max}$, active depth ℓ, and temperature T, the product $\Pi = \rho_t \eta_e$ of causal density and energy fidelity is bounded by $\kappa = p_{\max}/(\ell\, kT \ln 2)$, and architectures of equal product form genuine equivalence classes — the orbits of a one-parameter rescaling action — with

the mismatch penalty for a throttled architecture depending only on the class invariant. For silicon under explicitly stated assumptions ($p_{\max} = 100\ \mathrm{W\,cm^{-2}}$, $\ell = 1\ \mathrm{mm}$, $T = 300\ \mathrm{K}$), the bound is $\kappa = 3.483 \times 10^{29}\ \mathrm{m^{-3}\,s^{-1}}$; the causal-cone ceiling exceeds it by 8.607×10^{10}, so power, not relativity, binds; and two assumed architecture profiles occupy distinct classes at 3.0% of the bound and at the bound itself in principle. The construction is deliberately falsifiable: a sustainably operated platform exceeding its independently characterized κ by more than an order of magnitude, or saturated platforms on one substrate spanning more than an order of magnitude in product, would refute the invariant and localize the failure in the Landauer floor, the power-budget model, or the causal-cone assumption. The framework's value, if confirmed, is a substrate-architecture map computed from three physical constants in place of per-pair benchmarking; its principal limitation, inherited from its single-scalar form, is that it classifies thermodynamic fit and not computational adequacy.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.45.1-codeagent). 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: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 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] Equivalence Classes of Colorings. arXiv:1208.0993. https://arxiv.org/abs/1208.0993 [2] Equivalence Classes of Staged Trees. arXiv:1512.00209. https://arxiv.org/abs/1512.00209 [3] Independence equivalence classes of cycles. arXiv:2104.10080. https://arxiv.org/abs/2104.10080 [4] Defining Functions on Equivalence Classes. arXiv:1907.07591. https://arxiv.org/abs/1907.07591 [5] Equivalence classes for large deviations. arXiv:1208.5699. https://arxiv.org/abs/1208.5699 [6] Independence equivalence classes of paths. arXiv:2105.01592. https://arxiv.org/abs/2105.01592 [8] Graph Equivalence Classes for Spectral Projector-Based Graph Fourier Transforms. arXiv:1701.02864. https://arxiv.org/abs/1701.02864 [9] Enumerations of Cayley graphs. arXiv:math/0609462. https://arxiv.org/abs/math/0609462 [10] The Domination Equivalence Classes of Paths. arXiv:1710.03871. https://arxiv.org/abs/1710.03871 [12] https://papers.qnfo.org/papers/biological-and-engineered-computational-systems. QNFO: Biological and Engineered Computational Systems. [13] https://papers.qnfo.org/papers/computation-as-hamiltonian-isomorphism-a-mode-based-taxonomy-of-physical-com-

puti. QNFO: Computation as Hamiltonian Isomorphism: A Mode-Based Taxonomy of Physical Computing Substrates. [14] https://papers.qnfo.org/papers/comprehensive-technical-framework-for-network-isomorphism. QNFO: Comprehensive Technical Framework for Network Isomorphism.