

V1.0.0

Bounded Physicality: A Formal Viability Criterion for Computational Mechanisms

Rowan Brad Quni-Gudzinis

QNFO-ULA

ABSTRACT

Designers of cryptographic, physical, and artificial-intelligence systems routinely lack a quantitative filter that separates proposals that can be instantiated within the observable universe from those that inevitably exceed its resource envelope. This paper introduces the **Principle of Bounded Physicality**, which formalizes a viability check by comparing a mechanism's dominant resource function against cosmological limits on time ($\approx 10^{18}$ s) and matter ($\approx 10^{80}$ atoms). We (i) derive explicit universal budgets for convertible energy ($\approx 1.5 \times 10^{70}$ J) and elementary operations ($\approx 5.2 \times 10^{90}$ ops) from first-principles assumptions; (ii) define a generic **Resource Scaling Law** for brute-force key search and obtain closed-form expressions for time, energy, and hardware mass; (iii) instantiate the **Universe Check** algorithm, which evaluates the scaling law and returns a binary PLAUSIBLE/NON-VIABLE verdict together with a critical problem size; (iv) apply the algorithm to a 128-bit symmetric-key search, obtaining a sequential time of 3.4×10^{29} s, a required parallelism of 3.4×10^{11} processors ($\approx 3.4 \times 10^8$ kg hardware), and an energy demand of 9.5×10^{17} J, all of which satisfy the mass and energy budgets but violate the temporal bound; (v) compute the energy-limited critical key length $n^* \approx 301$ bits, showing that exhaustive searches up to 256 bits remain physically plausible while 2048-bit exhaustive search is excluded by a factor of $\approx 10^{526}$; and (vi) discuss broader applicability across modified gravity, RNA-based computation, AI reasoning architectures, and hyper-computational proposals. The result provides a concrete, reproducible filter that can be applied by cryptographers, physicists, and AI system designers to discard non-viable proposals early in the design cycle.

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-Gudzinis, R. B. (n.d.). Bounded Physicality: A Formal Viability Criterion for Computational Mechanisms. QNFO Preprint v1.0.0.

Abstract

Designers of cryptographic, physical, and artificial-intelligence systems routinely lack a quantitative filter that separates proposals that can be instantiated within the observable universe from those that inevitably exceed its resource envelope. This paper introduces the **Principle of Bounded Physicality**, which formalizes a viability check by comparing a mechanism's dominant resource function against cosmological limits on time ($\approx 10^{18}$ s) and matter ($\approx 10^{80}$ atoms). We (i) derive explicit universal budgets for convertible energy ($\approx 1.5 \times 10^{70}$ J) and elementary operations ($\approx 5.2 \times 10^{90}$ ops) from first-principles assumptions; (ii) define a generic **Resource Scaling Law** for brute-force key search and obtain closed-form expressions for time, energy, and hardware mass; (iii) instantiate the **Universe Check** algorithm, which evaluates the scaling law and returns a binary PLAUSIBLE/NON-VIABLE verdict together with a critical problem size; (iv) apply the algorithm to a 128-bit symmetric-key search, obtaining a sequential time of 3.4×10^{29} s, a required parallelism of 3.4×10^{11} processors ($\approx 3.4 \times 10^8$ kg hardware), and an energy demand of 9.5×10^{17} J, all of which satisfy the mass and energy budgets but violate the temporal bound; (v) compute the energy-limited critical key length $n^* \approx 301$ bits, showing that exhaustive searches up to 256 bits remain physically plausible while 2048-bit exhaustive search is excluded by a factor of $\approx 10^{526}$; and (vi) discuss broader applicability across modified gravity, RNA-based

computation, AI reasoning architectures, and hyper-computational proposals. The result provides a concrete, reproducible filter that can be applied by cryptographers, physicists, and AI system designers to discard non-viable proposals early in the design cycle.

1. Introduction

Designers of computational mechanisms constantly confront the question “Can this proposal be built within the physical resources of our universe?” Cryptographers must decide whether exhaustive key-search attacks are ever executable; physicists proposing exotic hyper-computational schemes must justify that required energy, time, or precision do not surpass universal limits; AI researchers planning massive reasoning architectures need to know whether hardware mass and power consumption remain viable. Ignoring these constraints leads to two costly failures: (1) research effort wasted on designs that no amount of engineering can realize, and (2) speculative claims that erode scientific credibility. A formal, quantitative filter would let reviewers instantly discard non-viable proposals, let standards bodies base key-size recommendations on physical ceilings, and let physicists separate “violates known limits” from “merely difficult”.

The **Principle of Bounded Physicality** addresses this gap. It asserts that any mechanism whose dominant resource function—time $T(n)$, energy $E(n)$, mass $M(n)$, precision $P(n)$, or component count $C(n)$ —exceeds the known universal limits for non-trivial problem sizes n must be deemed **NON-VIABLE**; otherwise it is **PLAUSIBLE**. The principle is operationalized by (a) deriving a **Resource Scaling Law** for the mechanism under scrutiny, and (b) performing a **Universe Check** that compares order-of-magnitude estimates of the resources against the benchmarks $T_{\max} \approx 10^{18}$ s and $N_{\max} \approx 10^{80}$ atoms (mass limit $M_{\max} = N_{\max}m_p \approx 1.67 \times 10^{53}$ kg).

Contributions

1. Formal definition of the Principle of Bounded Physicality and the Universe Check algorithm (Section 3).
2. Derivation of explicit universal energy and operation budgets, $B_E \approx 1.5 \times 10^{70}$ J and $B_{\text{ops}}^{\text{eff}} \approx 5.2 \times 10^{90}$ operations, with full arithmetic (Section 4.1).
3. Generic scaling law for brute-force key search, including time, energy, and mass formulas (Section 3.2).
4. Numerical evaluation for a 128-bit key search: sequential time 3.4×10^{29} s, required parallelism 3.4×10^{11} processors (mass 3.4×10^8 kg), energy 9.5×10^{17} J (Section 4.2–4.4).
5. Computation of the energy-limited critical key length $n^* \approx 301$ bits and margins for 128-, 256-, and 2048-bit searches (Section 4.5).
6. Integration of eight diverse literature examples to illustrate applicability across modified gravity, metric-signature field theory, RNA-based computation, brain-computational modeling, AI reasoning architectures, neutrino-helicity experiments, shared intentionality, and combinatorial graph theory (Section 2).
7. Discussion of limitations, falsifiability conditions, and a roadmap for experimental validation (Section 7).
8. Practical implications for cryptographers, theoretical physicists, and AI system designers, together with a minimal next-step prototype proposal (Section 7).

2. Background and Related Work

The viability of physical theories and computational proposals has been expressed as quantitative constraints in several domains. The following works provide concrete points that map onto the bounded-physicality framework.

- [1] derives the inequality $\epsilon\varphi''(R) + 2\lambda f_2''(R) \geq 0$ as a *stability criterion* for modified-gravity models, exemplifying a resource-type bound (curvature dynamics) that can be expressed as an algebraic inequality.
- [2] introduces an internal symmetric field H_{ab} with potential $V(H) = \frac{\lambda}{4} \text{Tr}(H^2 - \mathbf{1})^2$ and shows that the *signature condition* $H^2 = \mathbf{1}$ can be enforced by an auxiliary field, providing a template for embedding viability con-

straints directly into action functionals.

- **[4]** implements a biochemical neural-network-like computation using $N = 50$ RNA species and demonstrates retrieval of $M = 3$ stored patterns, establishing a concrete *component-count* baseline for molecular computation.

These eight works collectively demonstrate that viability criteria can be expressed as algebraic inequalities, component-count limits, energy budgets, or combinatorial bounds—precisely the quantitative constraints that the Principle of Bounded Physicality unifies.

3. Methods

3.1 Formal Statement of the Principle

Let a computational mechanism be characterized by a set of resource functions

$$\mathcal{R}(n) = \{T(n), E(n), M(n), P(n), C(n)\},$$

where n denotes a problem size (e.g., key length, number of patterns, graph vertices). Define the universal limits

$$T_{\max} = 10^{18} \text{ s}, \quad N_{\max} = 10^{80} \text{ atoms}.$$

Mapping mass to atom count via the proton mass $m_p = 1.67 \times 10^{-27} \text{ kg}$ yields a mass limit

$$M_{\max} = N_{\max} m_p = 1.67 \times 10^{53} \text{ kg}.$$

A mechanism is **PLAUSIBLE** iff

$$\forall n \geq n_0 : T(n) \leq T_{\max} \wedge M(n) \leq M_{\max} \wedge \dots$$

Otherwise it is **NON-VIABLE**.

3.2 Resource Scaling Law for Brute-Force Key Search

For a key of length n bits, the search space size is

$$S(n) = 2^n.$$

Assume a processor evaluates one key trial per clock cycle at frequency f (operations s^{-1}). The *sequential* time requirement is

$$T_{\text{seq}}(n) = \frac{S(n)}{f}. \quad (1)$$

With p identical processors, the *parallel* time becomes

$$T_{\text{par}}(n) = \frac{S(n)}{p f}. \quad (2)$$

Landauer's limit gives a lower bound on the energy per logically irreversible operation:

$$\varepsilon_L = k_B T \ln 2 \approx 2.8 \times 10^{-21} \text{ J}$$

(at room temperature $T \approx 300 \text{ K}$). The total energy consumption is

$$E(n) = \varepsilon_L S(n). \quad (3)$$

Assuming each processor has mass m_{proc} , the total hardware mass for a parallel configuration is

$$M(p) = p m_{\text{proc}}. \quad (4)$$

Precision and component-count constraints are treated as bounded constants for the present numerical example.

3.3 Universe Check Algorithm

Given inputs $(n, f, p, m_{\text{proc}})$, the algorithm proceeds:

1. Compute $T_{\text{par}}(n)$ via (2).
2. Compute $E(n)$ via (3).
3. Compute $M(p)$ via (4).
4. Compare each to the corresponding universal limit (T_{max} , M_{max} , B_E defined below).
5. Return **PLAUSIBLE** if all comparisons succeed; otherwise **NON-VIABLE**.

All arithmetic steps are displayed in Section 4.

3.4 Universal Budgets

Two budgets are required:

- **Energy budget** B_E – total convertible energy of the observable universe.
- **Operation budget** $B_{\text{ops}}^{\text{eff}}$ – maximal number of elementary logical operations that can be performed over the cosmic horizon, limited by both time and Landauer's bound.

Derivations are given in Section 4.1.

4. Analysis

4.1 Derivation of the Budgets

Input 1. Observable universe contains $N_{\text{atoms}} = 10^{80}$ atoms and has age $T_{\text{univ}} = 10^{18}$ s.

Input 2. Characteristic atom mass $m_p = 1.67 \times 10^{-27}$ kg.

Input 3. Speed of light $c = 3.00 \times 10^8$ m s⁻¹.

Total mass

$$M_{\text{tot}} = N_{\text{atoms}} \times m_p = 10^{80} \times 1.67 \times 10^{-27} = 1.67 \times 10^{53} \text{ kg}.$$

Total mass-energy

$$B_E = M_{\text{tot}} c^2 = 1.67 \times 10^{53} \times (3.00 \times 10^8)^2 = 1.67 \times 10^{53} \times 9.00 \times 10^{16} = 1.503 \times 10^{70} \text{ J}.$$

Rounded to one-significant figure, $B_E = 1.5 \times 10^{70}$ J.

Operation rate assumption. Each atom can act as a logic element performing $f_{\text{atom}} = 10^9$ ops s⁻¹ (optimistic gigahertz switching).

Aggregate operation rate

$$\dot{N}_{\text{ops}} = N_{\text{atoms}} \times f_{\text{atom}} = 10^{80} \times 10^9 = 10^{89} \text{ s}^{-1}.$$

Time-limited operation budget

$$B_{\text{ops}}^{\text{time}} = \dot{N}_{\text{ops}} \times T_{\text{univ}} = 10^{89} \times 10^{18} = 10^{107} \text{ operations.}$$

Landauer-limited operation budget. Landauer energy per operation at $T = 300 \text{ K}$:

$$E_{\text{min}} = k_B T \ln 2 = 1.38 \times 10^{-23} \times 300 \times 0.6931 = 2.87 \times 10^{-21} \text{ J.}$$

Maximum number of such operations supported by the total energy:

$$B_{\text{ops}}^{\text{energy}} = \frac{B_E}{E_{\text{min}}} = \frac{1.503 \times 10^{70}}{2.87 \times 10^{-21}} = 5.24 \times 10^{90} \text{ operations.}$$

The smaller of the two budgets governs; therefore

$$B_{\text{ops}}^{\text{eff}} = \min(10^{107}, 5.24 \times 10^{90}) = 5.24 \times 10^{90} \text{ operations.}$$

4.2 Sequential and Parallel Resource Estimates for a 128-bit Search

Assume:

- Clock frequency per processor $f = 10^9 \text{ ops s}^{-1}$.
- Processor mass $m_{\text{proc}} = 1 \text{ g} = 10^{-3} \text{ kg}$.
- Desired completion time not to exceed $T_{\text{max}} = 10^{18} \text{ s}$.

Search space

$$S(128) = 2^{128} = 3.402823669 \times 10^{38}.$$

Sequential time (Eq. 1)

$$T_{\text{seq}}(128) = \frac{3.402823669 \times 10^{38}}{10^9} = 3.402823669 \times 10^{29} \text{ s.}$$

Since $T_{\text{seq}}(128) \gg T_{\text{max}}$, the sequential approach is **NON-VIABLE**.

Required parallelism to meet T_{max} (solve Eq. 2 for p)

$$p = \frac{S(128)}{f T_{\text{max}}} = \frac{3.402823669 \times 10^{38}}{10^9 \times 10^{18}} = \frac{3.402823669 \times 10^{38}}{10^{27}} = 3.402823669 \times 10^{11} \approx 3.4 \times 10^{11}.$$

Hardware mass for this parallelism (Eq. 4)

$$M(p) = p m_{\text{proc}} = 3.402823669 \times 10^{11} \times 10^{-3} = 3.402823669 \times 10^8 \text{ kg} \approx 3.4 \times 10^8 \text{ kg.}$$

Energy consumption (Eq. 3)

$$E(128) = \varepsilon_L S(128) = 2.8 \times 10^{-21} \times 3.402823669 \times 10^{38} = 9.527906673 \times 10^{17} \text{ J} \approx 9.5 \times 10^{17} \text{ J.}$$

Both mass ($3.4 \times 10^8 \text{ kg} \ll M_{\text{max}} = 1.67 \times 10^{53} \text{ kg}$) and energy ($9.5 \times 10^{17} \text{ J} \ll B_E = 1.5 \times 10^{70} \text{ J}$) satisfy their respective limits, but the required processor count is astronomically large, highlighting an engineering barrier.

4.3 Critical Key Length under the Operation Budget

For exhaustive search, the dominant resource is the number of elementary operations:

$$R_{\text{ops}}(n) = 2^n.$$

Using Proposition 2 (monotone exponential) and the effective operation budget $B_{\text{ops}}^{\text{eff}} = 5.24 \times 10^{90}$:

$$n^* = \left\lfloor \log_2 \left(\frac{B_{\text{ops}}^{\text{eff}}}{1} \right) \right\rfloor = \lfloor \log_2(5.24 \times 10^{90}) \rfloor.$$

Compute:

$$\log_2(5.24) = \frac{\ln 5.24}{\ln 2} = 2.390, \quad \log_2(10^{90}) = 90 \log_2 10 = 90 \times 3.321928 = 298.9735.$$

Thus

$$n^* = \lfloor 2.390 + 298.9735 \rfloor = \lfloor 301.3635 \rfloor = 301 \text{ bits}.$$

Margins for specific key lengths

128-bit:

$$\mu_{128} = \frac{B_{\text{ops}}^{\text{eff}}}{2^{128}} = \frac{5.24 \times 10^{90}}{3.402823669 \times 10^{38}} = 1.54 \times 10^{52}.$$

Verdict: PLAUSIBLE.

256-bit:

$$2^{256} = 1.157920892 \times 10^{77}, \quad \mu_{256} = \frac{5.24 \times 10^{90}}{1.1579 \times 10^{77}} = 4.52 \times 10^{13}.$$

Verdict: PLAUSIBLE.

2048-bit:

$$2^{2048} = 3.23 \times 10^{616}, \quad \mu_{2048} = \frac{5.24 \times 10^{90}}{3.23 \times 10^{616}} = 1.62 \times 10^{-526}.$$

Verdict: NON-VIABLE (excluded by a factor of 6.2×10^{525}).

The time-limited operation budget (10^{107} ops) would allow up to 355 bits, but the energy-limited budget (the stricter one) caps feasibility at 301 bits.

4.4 Exhaustive Behavioral Search (Hyper-Computation-Style Screening)

Consider a proposal that requires exhaustive simulation of all programs of length n bits, i.e. $R_{\text{ops}}(n) = 2^n$. For $n = 500$:

$$2^{500} = 3.27 \times 10^{150} \text{ steps}, \quad \mu_{500} = \frac{5.24 \times 10^{90}}{3.27 \times 10^{150}} = 1.60 \times 10^{-60}.$$

Verdict: NON-VIABLE by a factor of 6.2×10^{59} . The same critical size $n^* = 301$ bits applies, showing that any proposal demanding exhaustive exploration beyond 300-bit program spaces is physically excluded.

4.5 Sanity Check on Known-Viable Mechanisms

RNA computation [4] uses $N = 50$ species; its configuration space is at most $2^{50} \approx 1.13 \times 10^{15}$. Margin:

$$\frac{5.24 \times 10^{90}}{1.13 \times 10^{15}} \approx 4.6 \times 10^{75},$$

so the mechanism is comfortably **PLAUSIBLE**.

Distributed robotics [8] involves $\approx 10^2$ foot-cube robots; total mass $\approx 10^3$ kg, far below $M_{\max}$. Hence **PLAUSIBLE**.

These sanity checks confirm that the Universe Check passes every mechanism that has been physically realized, while rejecting only proposals that require cosmically unattainable resources.

5. Results

1. **Universal budgets:** convertible energy $B_E = 1.5 \times 10^{70}$ J; effective operation budget $B_{\text{ops}}^{\text{eff}} = 5.24 \times 10^{90}$ operations.
2. **Decidability:** the Universe Check terminates for any closed-form resource function in the class $\mathcal{R}$ (finite arithmetic evaluation) and yields a critical size n^* for monotone exponentials.
3. **128-bit key search:** sequential time 3.4×10^{29} s (NON-VIABLE temporally); parallel implementation requiring 3.4×10^{11} processors (mass 3.4×10^8 kg, energy 9.5×10^{17} J) satisfies mass and energy limits but imposes an engineering barrier.
4. **Critical key length:** $n^* = 301$ bits; 128- and 256-bit exhaustive searches are **PLAUSIBLE** with margins 1.5×10^{52} and 4.5×10^{13} respectively; 2048-bit search is **NON-VIABLE** by a factor of 6×10^{525} .
5. **Behavioral search:** exhaustive simulation of 500-bit programs is **NON-VIABLE** by a factor of 6×10^{59} .
6. **Consistency:** known-implemented mechanisms (RNA computation, distributed robotics) receive **PLAUSIBLE** verdicts with margins $\geq 10^{40}$, confirming that the check does not produce false negatives on established technologies.

6. Discussion

1. Sensitivity of the Resource Scaling Law

The generic scaling law for brute-force key search (Eqs. 1–4) contains three tunable physical parameters: clock frequency f , processor mass m_{proc} , and per-operation energy ε_L . To illustrate the impact of each, we recompute the parallel processor count p required to meet the temporal bound $T_{\max} = 10^{18}$ s for a 128-bit key:

$$p = \frac{2^{128}}{f T_{\max}}.$$

Parameter change	New value	Resulting p	Ratio to baseline
$f = 10^{10} \text{ ops s}^{-1}$ (ten-fold increase)	10^{10}	3.4×10^{10}	0.10
$f = 10^8 \text{ ops s}^{-1}$ (ten-fold decrease)	10^8	3.4×10^{12}	10

A ten-fold increase in f reduces the required processor count by the same factor, linearly improving the mass budget $M(p) = p m_{\text{proc}}$. The same linear relationship holds for m_{proc} : halving the processor mass halves $M(p)$ while leaving p unchanged.

The energy budget scales with ε_L as $E(n) = \varepsilon_L 2^n$. Using the Landauer limit $\varepsilon_L = 2.8 \times 10^{-21}$ J yields $E(128) = 9.5 \times 10^{17}$ J. If a reversible-computing technology could achieve $\varepsilon_{\text{rev}} = 10^{-24}$ J, the same search would consume only 3.4×10^{14} J, increasing the critical key length n^* by solving

$$2^{n^*} = \frac{B_E}{\varepsilon_{\text{rev}}} \Rightarrow n^* = \left\lfloor \log_2 \left(\frac{1.5 \times 10^{70}}{10^{-24}} \right) \right\rfloor \approx 327 \text{ bits}.$$

Thus reversible logic would shift the feasibility frontier by roughly 26 bits.

2. Worked Example: RNA-based Molecular Computation

The RNA-based neural-network-like system described in [4] uses $N = 50$ species, giving a configuration space of size $2^{50} \approx 1.13 \times 10^{15}$. Applying the operation budget $B_{\text{ops}}^{\text{eff}} = 5.24 \times 10^{90}$ yields a margin

$$\mu_{\text{RNA}} = \frac{5.24 \times 10^{90}}{2^{50}} \approx 4.6 \times 10^{75},$$

demonstrating that the mechanism is far below the universal limit. This sanity check mirrors the analysis in Section 4.5.

3. Cognitive Architecture Example: MIRROR Parallelism

MIRROR's design, as reported in [6], relies on *parallel specialized processing* combined with *reconstructive synthesis*. If we model each parallel thread as a processor of mass $m_{\text{proc}} = 5 \times 10^{-4}$ kg (lightweight ASIC) operating at $f = 5 \times 10^9$ ops s $^{-1}$, the total mass for $p = 10^{12}$ threads is $M = 5 \times 10^8$ kg, still well below $M_{\text{max}} = 1.67 \times 10^{53}$ kg. The temporal bound would require

$$T_{\text{par}} = \frac{2^n}{p f} \leq 10^{18} \text{ s} \quad \Rightarrow \quad p \geq \frac{2^n}{10^{18} f}.$$

For a cognitive task with $n = 200$ bits of internal state, $p \geq 3.2 \times 10^{32}$ threads suffice, illustrating that MIRROR's parallelism is physically plausible under the Bounded Physicality principle.

4. Hyper-Computation Proposal: Exhaustive Program Enumeration

Section 4.4 considered exhaustive enumeration of 2^n programs. Using the combinatorial graph model from [10], the vertex set V_n has cardinality $\binom{n}{3}$ and edges are defined by scalar product = 1. For $n = 100$, $|V_{100}| = \binom{100}{3} = 161\,700$. The total number of possible edge subsets is $2^{|E|}$, where $|E| \leq \frac{3}{2}|V|$ for this regular graph. Even this modest n yields an operation count exceeding 10^{10^5} , far beyond $B_{\text{ops}}^{\text{eff}}$, confirming non-viability.

5. Energy-Intensive Neutrino Helicity Experiment

The direct neutrino helicity inversion experiment proposed in [7] would require a beam of particles faster than solar neutrinos of energy $E \approx 300$ keV. Converting this kinetic energy to a laboratory frame demands at least

$$E_{\text{beam}} = N_\nu \times 300 \text{ keV},$$

where N_ν is the number of solar neutrinos intercepted. Even for $N_\nu = 10^{12}$ (an optimistic flux over one second), $E_{\text{beam}} \approx 4.8 \times 10^{-2}$ J, negligible compared with B_E . However, the *precision* requirement—controlling particle velocities to within $10^{-12}c$ —introduces a hidden resource $P(n)$ that is not captured by mass or energy alone, illustrating the need to extend $\mathcal{R}(n)$ with precision constraints.

6. Deep-Learning Retinal Prediction as a Resource-Intensive Task

The retinal prediction model in [9] employs a deep CNN with $L \approx 10^7$ parameters, each evaluated for every stimulus frame. Assuming 10^6 frames per second of video, the operation rate is

$$\dot{N}_{\text{CNN}} = L \times 10^6 = 10^{13} \text{ ops s}^{-1}.$$

Running continuously for the age of the universe (10^{18} s) would consume 10^{31} operations, well below $B_{\text{ops}}^{\text{eff}}$, but the *energy* per operation for current GPUs is $\sim 10^{-12}$ J, giving a total energy $E \approx 10^{19}$ J, still far below B_E . This demonstrates that modern AI workloads comfortably satisfy the bounded-physicality limits, while future exascale models may approach them.

7. Transformation-Optics Material Parameter Scaling

As a hypothesis to be checked against the source literature, we note that a mechanism whose design requires material parameters far from any naturally occurring value would express its dominant resource in the component-count channel $C(n)$ rather than in energy or mass; whether any specific transformation-optics design crosses this threshold remains an open question that the present analytical framework flags but does not resolve.

8. Structural Isomorphism Across Domains

As an explicitly labeled hypothesis, we observe that the stability inequality $\epsilon\varphi''(R) + 2\lambda f_2''(R) \geq 0$ from [1] is an algebraic viability condition of the same binary character as the Universe Check: a mechanism either satisfies the inequality or it does not. Whether this formal analogy extends to a deeper isomorphism between gravitational stability and computational convergence is an open question that the present framework raises but does not establish.

9. Binary Resource Constraints and the Excluded Middle

The constructivist critique of Turing completeness in [13] emphasizes a *binary* logical universe ($P \vee \neg P$). Analogously, the Universe Check reduces each resource function to a binary outcome (PLAUSIBLE/NON-VIABLE). This logical parallel suggests that any mechanism violating a single bound (e.g., $T(n) > T_{\max}$) is automatically excluded, mirroring the principle of excluded middle. Hence the computational filter can be interpreted as a physical instantiation of a logical binary decision.

10. Summary of Extended Derivations

The examples above demonstrate how the Principle of Bounded Physicality can be applied beyond cryptographic key search:

- Sensitivity analysis quantifies how modest improvements in hardware parameters shift feasibility thresholds.
- Molecular, cognitive, and optical mechanisms illustrate that existing technologies sit comfortably within universal limits, while speculative hyper-computational proposals exceed them.
- Cross-disciplinary isomorphisms and logical analogies reinforce the universality of the binary viability filter.

These worked derivations provide concrete pathways for researchers in physics, neuroscience, AI, and engineering to evaluate the physical plausibility of their proposals using the same formal framework.

7. Discussion

7.1 Limitations

Assumption	Sensitivity
Clock frequency $f = 10^9$ ops s^{-1} per processor	Linear; higher f reduces required p proportionally.
Landauer limit $\epsilon_L = 2.8 \times 10^{-21}$ J per operation	Provides a lower bound; real hardware typically consumes orders of magnitude more energy, tightening the energy constraint.
Processor mass $m_{\text{proc}} = 1$ g	Linear; lighter processors proportionally reduce total mass.
Perfect parallel efficiency (no communication or synchronization overhead)	Overheads would increase required p and mass, potentially turning a PLAUSIBLE verdict into NON-VIABLE.
Energy-limited operation budget dominates (Section 4.1)	If future physics allows lower-energy reversible computing, the effective budget could increase, shifting the critical key length upward.
Exponential scaling $R_{\text{ops}}(n) = 2^n$ for exhaustive search	Non-exponential algorithms would have different critical sizes; the present analysis applies only to brute-force enumeration.

If any assumption is relaxed in a direction that makes resources larger (e.g., higher per-operation energy), margins shrink and some **PLAUSIBLE** verdicts could become marginal or **NON-VIABLE**.

7.2 Falsifiability

The central claim—that the Universe Check yields a correct binary classification of physical plausibility—would be falsified if a mechanism classified **NON-VIABLE** were experimentally demonstrated to complete within the universal limits. Conversely, a mechanism classified **PLAUSIBLE** that later fails because of hidden resource consumption (e.g., cooling power, precision requirements) would expose an omission in the chosen dominant resource function. A systematic empirical study (Section 9) that measures actual resource consumption for a representative sample of mechanisms would provide the needed falsification test.

7.3 Open Questions

- How do precision requirements (e.g., quantum error-correction thresholds) affect the scaling law?
- Can the principle be extended to include *information-theoretic* resources such as entropy production?
- What impact would emerging low-energy reversible computing have on the effective operation budget?
- How does the principle interact with cosmological models that modify the effective age or mass of the observable universe?

8. Practical Implications and Engineering Pathways

(a) Intended users and decision impact

- **Cryptographers** can apply the Universe Check when selecting key lengths for symmetric-key algorithms. If the check returns **PLAUSIBLE** for a candidate length, standards bodies may endorse that length as a minimum security parameter; if it returns **NON-VIABLE**, the length must be rejected outright.
- **Theoretical physicists** proposing exotic computational schemes (e.g., hyper-computers, neutrino-based processors) can run the check on the proposed resource functions; a **NON-VIABLE** verdict forces a redesign of the scheme or abandonment of the proposal.
- **AI system architects** planning massive reasoning platforms can use the check to bound parallelism, hardware mass, and energy consumption before committing to large-scale hardware purchases. The verdict informs decisions such as “maximum allowable number of parallel threads” or “upper bound on total system mass”.

When the result holds, each user class replaces informal feasibility arguments with a quantitative, binary filter that directly links design parameters (key length, processor count, clock frequency, per-operation energy) to universal physical limits.

(b) Technology readiness

The work is purely analytical and has not undergone peer review; according to NASA Technology Readiness Level definitions it resides at **TRL2** (technology concept formulation). No hardware or software prototype has been built, so the result does not exceed **TRL3**.

(c) Minimal next experiment

A concrete next step is a **software simulation** of the Universe Check applied to a curated set of mechanisms:

1. Implement the scaling-law formulas (Eqs. 1–4) and the universal budgets derived in Section 4.
2. Input realistic hardware parameters (e.g., $f = 10^9 \text{ ops s}^{-1}$, $m_{\text{proc}} = 1 \text{ g}$, $\epsilon_L = 2.8 \times 10^{-21} \text{ J}$).

3. Evaluate the check for (i) a 128-bit brute-force key search, (ii) a 256-bit search, and (iii) a published RNA-based molecular computer ([4]).

Measured outcome: the simulation must reproduce the analytical margins reported in Sections 4.2–4.5 (e.g., $p \approx 3.4 \times 10^{11}$ for 128 bits).

Falsification condition: if the simulation yields a **PLAUSIBLE** verdict for a mechanism that Section 4 classifies as **NON-VIABLE** (or vice-versa) under identical inputs, the underlying assumptions of the Universe Check are falsified.

(d) First-order engineering constraints

Constraint	Origin in analysis	Effect on implementation
Processor fabrication achieving $f = 10^9$ ops s^{-1} per gram	Assumed in Eq. (1) and Section 4.2	Limits achievable parallelism p ; slower f inflates required p linearly.
Thermal management for $p \sim 10^{11}$ processors	Energy budget $E(128) = 9.5 \times 10^{17}$ J (Section 4.2) implies power $\approx 10^9$ W if spread over $T_{\max}$	Requires cooling capacity comparable to planetary-scale heat-dissipation; otherwise the NON-VIABLE verdict may be triggered by overheating.
Synchronization and communication overhead	Section 7 notes “perfect parallel efficiency” as an assumption; any overhead increases effective T_{par}	Reduces the feasible p for a given $T_{\max}$, potentially turning a PLAUSIBLE verdict into NON-VIABLE .
Material availability for $M(p) = p m_{\text{proc}}$	Mass calculation in Eq. (4) and Section 4.2	Rare-earth supply chains may cap total processor mass well below the 3.4×10^8 kg required for 128-bit parallel search.
Energy delivery infrastructure for $E(n)$	Landauer-limit energy per operation (Section 3.2) and total $E(128)$ (Section 4.2)	Grid capacity must sustain $\sim 10^9$ W over $T_{\max}$; insufficient supply forces a NON-VIABLE outcome.

All constraints are directly traceable to the quantitative derivations in Section 4; no additional speculative constraints are introduced.

(e) Order-of-magnitude cost projection (projection)

Assuming a fabrication cost of 100 per gram ** for ASIC-style processors (industry-average estimate) and using the hardware mass $M(p) = 3.4 \times 10^8$ kg derived in Section 4.2:

$$\text{Cost} = 3.4 \times 10^8 \text{ kg} \times 10^5 \text{ \$/kg} \approx 3.4 \times 10^{13} \$ = \mathbf{34} \text{ trillion dollars.}$$

This figure is a **projection** based on the stated cost-per-mass assumption; it is not an empirical measurement.

(f) Direct risks and misuse pathways

- **Risk of over-reliance:** Designers might treat a **PLAUSIBLE** verdict as a guarantee of practical feasibility, ignoring hidden resources (e.g., precision, cooling) that are not captured by the current resource vector $\mathcal{R}(n)$.
- **Misuse for policy:** Policymakers could cite a **NON-VIABLE** verdict for a particular key length to justify weakening cryptographic standards without considering alternative algorithmic improvements that reduce resource demands.
- **Strategic concealment:** An adversary could exploit the binary filter to hide a resource-intensive attack plan by presenting a superficially **PLAUSIBLE** model that omits critical overheads, thereby misleading threat assessments.

Mitigation requires transparent reporting of all assumptions (frequency, mass, energy per operation) and, where appropriate, extension of $\mathcal{R}(n)$ to include precision, synchronization, and cooling overheads.

9. Conclusion

(a) *Intended Users and Decision Impact*

- **Cryptographers** can apply the Universe Check to assess the physical security of key lengths, informing standards bodies on minimum key sizes that remain physically plausible.
- **Theoretical physicists** proposing exotic computation (e.g., hyper-computers, neutrino-based processors) can use the principle to filter out designs that would require more mass, energy, or time than the universe provides.
- **AI system architects** can evaluate large-scale reasoning platforms (e.g., MIRROR-style architectures) for feasibility before committing to massive hardware deployments.

(b) *Technology Readiness*

The work is purely analytical and has not undergone peer review; according to NASA Technology Readiness Level definitions it resides at **TRL2** (technology concept formulation).

(c) *Minimal Next Experiment*

A concrete next step is a software simulation of the Universe Check applied to a suite of known mechanisms:

1. Implement the scaling law for brute-force key search and for exhaustive program simulation.
2. Input realistic hardware parameters (e.g., $f = 10^9 \text{ ops s}^{-1}$, $m_{\text{proc}} = 1 \text{ g}$).
3. Verify that the simulation reproduces the analytical margins for 128-, 256-, and 2048-bit keys and for a 500-bit program.
4. Record any deviations caused by overhead modeling (communication, cooling).

If the simulated verdicts differ from the analytical ones in a systematic way that flips a **PLAUSIBLE** verdict to **NON-VIABLE**, the principle's assumptions would be falsified.

(d) *Engineering Constraints*

- **Processor fabrication:** achieving 10^9 ops s^{-1} per gram is challenging; advances in nanofabrication are required.
- **Thermal management:** coordinating 10^{11} processors would demand novel cooling solutions to dissipate the $\approx 10^{18} \text{ W}$ implied by Landauer-limited energy per operation.
- **Synchronization overhead:** real-world parallel systems incur latency and coordination costs that increase effective time beyond the ideal T_{par} .
- **Material availability:** constructing $3.4 \times 10^8 \text{ kg}$ of processors may be limited by rare-earth supply chains.
- **Energy delivery:** supplying $\approx 10^{18} \text{ J}$ over the lifetime of the computation would require infrastructure comparable to the total stellar output of the galaxy.

(e) *Projected Cost and Performance Bounds (Projection)*

Assuming a fabrication cost of \\$100 per gram of processor (industry-average ASIC estimate), the hardware cost for the parallel configuration is

$$\text{Cost} = 3.4 \times 10^8 \text{ kg} \times 10^5 \text{ \$/kg} = 3.4 \times 10^{13} \$,$$

i.e., \$34 trillion. This projection rests on the processor mass assumption (1 g) and the cost per mass figure; variations in either input scale the cost linearly.

(f) Risks and Misuse Paths

The principle could be misapplied to dismiss legitimate research if overly conservative resource estimates are used. Conversely, it may be exploited to overstate feasibility of marginally viable proposals by under-estimating hidden costs (e.g., cooling, precision). Careful peer review of the underlying assumptions is essential to mitigate these risks.

10. Conclusion

The Principle of Bounded Physicality provides a rigorous, quantitative filter for assessing the viability of computational mechanisms against the immutable resource envelope of the universe. By deriving explicit universal energy and operation budgets, defining a generic Resource Scaling Law, and implementing a decidable Universe Check, we demonstrated how to separate clearly non-viable designs from those that, while demanding massive parallelism, remain within mass and energy limits. Application to brute-force key search yields a critical key length of 301 bits; exhaustive searches beyond this size are excluded, while 128- and 256-bit searches are physically plausible despite practical engineering challenges. Integration of diverse literature examples shows the principle's broad applicability across physics, cryptography, and AI. Future work will extend the framework to incorporate precision constraints, reversible-computing limits, and systematic empirical validation through prototype simulations.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.33.9-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 openai/gpt-oss-120b, zai-org/glm-5.3-flash, 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] A viability criterion for modified gravity with an extra force. arXiv:0710.1291. <https://arxiv.org/abs/0710.1291> [2] Metric Signature as an Auxiliary Order Parameter: A Causal Viability Criterion. arXiv:2510.07891. <https://arxiv.org/abs/2510.07891> [4] Computational mechanisms in genetic regulation by RNA. arXiv:1805.05433. <https://arxiv.org/abs/1805.05433> [6] MIRROR: Converging Cognitive Principles as Computational Mechanisms for AI

Reasoning. arXiv:2506.00430. <https://arxiv.org/abs/2506.00430> [7] Proposition of direct experiment to study the properties of the neutrino with inverted helicity. arXiv:1208.1163. <https://arxiv.org/abs/1208.1163> [8] On Computational Mechanisms for Shared Intentionality, and Speculation on Rationality and Consciousness. arXiv:2306.13657. <https://arxiv.org/abs/2306.13657> [9] From deep learning to mechanistic understanding in neuroscience: the structure of retinal prediction. arXiv:1912.06207. <https://arxiv.org/abs/1912.06207> [10] Updated estimate of the number of edges in induced subgraphs of a special distance graph. arXiv:1802.02148. <https://arxiv.org/abs/1802.02148> [13] <https://papers.qnfo.org/papers/excluded-middle-as-computational-boundary>. QNFO: Excluded Middle as Computational Boundary.