V1.0.0

Constant-Gap Sculpted Hamiltonians for Factoring and NP-Hard Optimization: Analytic Gap Scaling and the Location of the Exponential Cost

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Analog Hamiltonian devices are widely dismissed because encoding unstructured computational problems is believed to demand control precision that grows exponentially with problem size. A proposed escape is to sculpt the Hamiltonian so that every invalid configuration is lifted by a constant energy gap ΔE, making correctness a matter of energetic protection ($k_B T \ll \Delta E$) rather than of analog precision. Whether such a constant-gap construction genuinely avoids exponential resources has not been settled by a single explicit derivation. We analyze a sculpted factoring Hamiltonian H_{fac} that combines a logarithmic product constraint with integer-valued penalty terms enforcing integrality and primality. Under the discrete-penalty convention, in which the product penalty $|pq - N|^2$ carries unit coefficient, we derive a ground-state energy $E_0 = 0$, a first excited state at $E_1 = 1$, and a constant gap $\Delta E = 1$ independent of the integer N; the corresponding measurement time at the experimentally reported spectroscopic transition rate of $70\ \mathrm{s}^{-1}$ is $T_{\text{meas}} \approx 1.43 \times 10^{-2}$ s, negligible against the $4{,}500$ h aggregate budget reported for a long-horizon analog design benchmark. Under the alternative smooth-encoding convention, in which only the logarithmic constraint is engineered, we show that the minimal non-zero constraint value scales as $1/N^2$, so that maintaining a thermal-protection gap at 10 mK for a 128-bit integer requires a coefficient $\alpha \approx 1.6 \times 10^{52}$ J. The exponential cost of analog factoring therefore does not vanish; it migrates from parameter precision to the physical implementation of arithmetic penalties. This reframes analog computation as a compilation-style engineering problem whose binding constraint is identifiable and testable.

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.). Constant-Gap Sculpted Hamiltonians for Factoring and NP-Hard Optimization: Analytic Gap Scaling and the Location of the Exponential Cost. QNFO Preprint v1.0.0.

1. Introduction

Factoring large integers underpins modern public-key cryptography, and the inability of classical computers to factor n-bit RSA moduli in polynomial time secures digital transactions at global scale. Engineers who attempt to embed factoring, or any NP-hard constraint-satisfaction problem, into a physical energy landscape face a concrete obstacle: the control parameters of an analog processor must distinguish the correct solution from exponentially many near-degenerate configurations, and the prevailing assumption is that this demands precision growing as $O(2^n)$ in the bit-length n. If that assumption holds, fabrication tolerances and thermal stability become unattainable at cryptographically relevant sizes, and investment in analog hardware for cryptanalysis and large-scale optimization is blocked at the design stage.

A proposed resolution is Hamiltonian sculpting: encode the problem constraints explicitly so that the low-energy manifold contains only valid solutions, separated from all invalid configurations by a constant spectral gap ΔE. Under this proposal, correctness follows from energetic protection — thermal fluctuations of scale $k_B T$ must be far smaller than ΔE — and a digital readout of the final state yields the answer. The central question is whether the constant-gap claim

survives explicit analysis, or whether it hides an exponential cost in the coefficients of the penalty terms, in the precision with which those coefficients must be realized, or in the physical implementation of the arithmetic that the penalties encode.

We answer this question by deriving the gap scaling of a sculpted factoring Hamiltonian under two explicitly stated conventions, quantifying the thermodynamic requirements with experimentally grounded transition rates and budget figures, and locating precisely where any exponential burden reappears. Our contributions are:

1. **Gap derivation** — We prove that the sculpted factoring Hamiltonian with unit-coefficient integer-valued penalties has ground-state energy $E_0 = 0$ and first excited energy $E_1 = 1$, giving a constant gap $\Delta E = 1$ independent of N (Section 4.1).
2. **Thermal and temporal analysis** — Using the transition rate $\Gamma = 70\ \mathrm{s}^{-1}$ reported for two-photon spectroscopy of molecular hydrogen ions [2], we compute the measurement time $T_{\mathrm{meas}} = 1/\Gamma \approx 1.43 \times 10^{-2}$ s and compare it with the $4{,}500$ h aggregate wall-clock budget reported for the analog design benchmark of [1] (Sections 4.2 and 4.3).
3. **Hidden-cost quantification** — Under the smooth-encoding convention we derive that the minimal non-zero value of the logarithmic constraint scales as $C_{\min} = 1/N^2$, and we compute that enforcing a thermal-protection gap $k_B T$ at $T = 10$ mK for a 128-bit integer requires a coefficient $\alpha \approx 1.6 \times 10^{52}$ J, forty orders of magnitude above the thermodynamic Protection Deficit of order 10^{12} reported for actively error-corrected quantum computation [7] (Section 4.4).
4. **Practical pathway** — We identify the users, the readiness level (TRL 1–2), the smallest falsifying prototype experiment, and the engineering constraints that bind first (Section 7).

Beyond quantum cryptanalysis, the result bears on any domain that maps NP-hard constraint-satisfaction problems onto physical energy landscapes, including combinatorial optimization hardware and adiabatic quantum processing; the analysis of where the exponential burden resides applies wherever smooth encodings of discrete constraints are considered.

2. Background and Related Work

The literature on analog design effort, precision spectroscopy, Hamiltonian simulation, and thermodynamic limits supplies the quantitative anchors for our analysis.

The long-horizon analog design benchmark of [1] evaluates agents on 50 transistor-level tasks across 15 configurations and 3 rollouts each, for a total of $2{,}250$ attempts at 2 h per attempt, and reports that current agents cannot yet guarantee functional analog circuits across operating conditions. We use the resulting $4{,}500$ h aggregate figure as the empirical scale of effort that the analog design community already accepts for exploring large design spaces, against which the cost of resolving a sculpted gap can be compared.

The precision spectroscopy of molecular hydrogen ions in [2] reports a two-photon transition rate of approximately $70\ \mathrm{s}^{-1}$ under a laser flux of $15\ \mathrm{W/mm}^2$ with an instrumental width of 10 kHz, and achieves relative accuracies of order 10^{-9} (1 ppb). This provides an experimentally validated rate at which an isolated quantum system can resolve a well-defined transition, which we adopt as the operational timescale for interrogating a sculpted energy gap.

The directional-statistics work of [3] analyzes estimators for the offset exponential and normal distributions in three-dimensional space and derives the explicit integral $\int_0^\infty r^2 e^{-r/l}\, dr = 2l^3$, a normalization result for radially decaying weight functions. We use this style of explicit integral normalization as a model for how penalty-term weightings in a sculpted Hamiltonian must be normalized so that no term silently changes the energy scale.

The fast-forwarding analysis of [4] examines the time-energy uncertainty principle (TEUP) and constructs counterexamples in which increasing the norm of the Hamiltonian shortens the measurement duration required to resolve an energy difference. This trade between Hamiltonian norm and measurement time is directly relevant here: it shows that the TEUP bound $\Delta t\,\Delta E \geq \frac{1}{2}$ is not a hard limit, and that pushing resolution into short times reappears as a demand on the engineered energy scale — exactly the migration of cost that our smooth-encoding analysis quantifies.

The sparse-Hamiltonian simulation result of [5] shows that a sparse Hamiltonian can be simulated with query complexity $O\big(\tau \frac{\log(\tau/\varepsilon)}{\log\log(\tau/\varepsilon)}\big)$, where $\tau = t\|H\|$ and ε is the simulation error, achieving an exponential improvement in precision for sparse Hamiltonians. Because the runtime retains the Hamiltonian norm $\|H\|$, this result cuts both ways for our problem: a bounded-norm sculpted Hamiltonian is efficiently simulable, while a construction whose norm grows as N^2 inherits that growth in any simulation or physical realization.

The construction of fermionic Hamiltonians with non-Gaussian low-energy states in [6] demonstrates that low-energy states with computationally useful structure can be engineered in specialized models, while relying on intricate mappings. This supports the plausibility of sculpted low-energy manifolds while cautioning that the mapping itself can carry resource costs — the same location-of-cost question our two conventions separate.

The thermodynamic analysis of [7] quantifies a "Protection Deficit" of order 10^{12} (dimensionless) in energy consumption for actively error-corrected quantum computation, establishing a thermodynamic wall for that architecture. Our smooth-encoding coefficient $\alpha \approx 1.6 \times 10^{52}$ J exceeds this deficit by forty orders of magnitude, which calibrates how far the constant-gap-by-large-coefficients route sits from any known thermodynamic budget.

The Programming-the-Vacuum framework of [8] treats Hamiltonian evolution itself as computation and distinguishes Class A (Ising-type optimization) from Class B (spectral synthesis), with synthetic performance generation demonstrated for problem sizes up to $N = 12$ spins. Our sculpted factoring Hamiltonian belongs to Class A under this taxonomy, and the small demonstrated sizes underline that scaling behavior, not small-instance function, is the open question.

Finally, the topological hydrodynamics study of [10] links factorization to spectral properties of superfluid vacua and concludes that for cryptographic integers ($N > 10^{50}$) the required energy gap falls far below any feasible thermal noise floor. This independent conclusion is consistent with the smooth-encoding scaling we derive in Section 4.4, and it frames the discrete-penalty route as the remaining candidate for evading the thermal-noise objection.

Collectively, these works establish that substantial analog design effort is already routine [1], that transition rates near $70\ \mathrm{s}^{-1}$ are experimentally real [2], that Hamiltonian-norm and measurement-time budgets trade against each other [4], [5], that thermodynamic walls of known magnitude constrain active correction [7], and that small-instance sculpted landscapes exist while their scaling remains the open issue [6], [8], [10].

3. Methods

Our analysis is analytical throughout; no numerical simulation, hardware experiment, or measurement is reported. It proceeds in three stages.

Stage 1: Hamiltonian construction. We define a sculpted Hamiltonian for integer factoring:

$$H_{\text{fac}}(p,q) = \big(\ln p + \ln q - \ln N\big)^2 + \lambda_1 \left|p\,q - N\right|^2 + \lambda_2 \left(\text{is_prime}(p) - 1\right)^2 + \lambda_3 \left(\text{is_prime}(q) - 1\right)^2,$$

where $p, q \in \mathbb{Z}^+$ with $2 \leq p \leq q \leq N$, $\text{is_prime}(\cdot)$ returns 1 for prime arguments and 0 otherwise, and the coefficients λ_i are the engineered energy scales of the penalty terms.

Stage 2: Convention selection. Two analyses of this Hamiltonian are possible, and they differ in which term carries the spectral gap. Under the **discrete-penalty convention**, the coefficients λ_i are set to 1 energy unit and the integer-valued product penalty $|pq - N|^2$ fixes the gap, because $|pq - N|$ is a non-negative integer for integer p, q, N. Under the **smooth-encoding convention**, only the logarithmic constraint $C(p, q; N) = (\ln p + \ln q - \ln N)^2$ is engineered, its coefficient α is a free physical energy scale, and the gap is $\Delta E = \alpha\, C_{\min}$. We adopt the discrete-penalty convention for the main gap result (Section 4.1) and treat the smooth-encoding coefficient as the quantification of where the exponential burden reappears when the discrete arithmetic penalty is not physically available (Section 4.4). This choice rests on the assumption that the product penalty can be implemented with unit coefficient and N-independent precision; Section 6 examines how sensitive the conclusion is to that assumption.

Stage 3: Thermal and temporal feasibility. Energetic protection requires $k_B T \ll \Delta E$. We take the transition rate $\Gamma = 70\ \mathrm{s}^{-1}$ from the two-photon spectroscopy of molecular hydrogen ions [2] as the proxy rate at which the system can resolve an energy difference of order ΔE, so that the characteristic measurement time is $T_{\text{meas}} = 1/\Gamma$. We take $k_B = 1.380649 \times 10^{-23}\ \mathrm{J\,K}^{-1}$ as a standard physical constant (stated here as our own assumption, as no source supplies it), and we assume thermal equilibrium with quantum tunneling neglected. For the budget comparison we use the analog design benchmark figures of [1]: 50 tasks, 15 configurations per task, 3 rollouts per configuration, and 2 h per attempt.

All arithmetic steps are displayed explicitly in Section 4.

4. Analysis

4.1 Gap derivation under the discrete-penalty convention

With $\lambda_1 = \lambda_2 = \lambda_3 = 1$ energy unit, we evaluate H_{fac} for three exhaustive cases:

Case	Condition	Energy
(i)	$pq = N$ and p, q prime	$\ln p + \ln q - \ln N = 0$, $\lvert pq - N\rvert^2 = 0$, both primality penalties 0; $E = 0$.
(ii)	$pq \neq N$, p, q prime	$\lvert pq - N\rvert \geq 1$ since pq and N are distinct integers, so $\lvert pq - N\rvert^2 \geq 1$; $E \geq 1$.
(iii)	p or q composite	At least one primality penalty contributes $(1-1)^2 = 0$ or $(0-1)^2 = 1$; for composite argument the penalty is exactly 1; $E \geq 1$.

Every invalid pair falls into case (ii) or (iii) or both, and in each the energy is at least 1 because each penalty term is a squared non-negative integer with unit coefficient, whose minimal non-zero value is 1. The unique ground state is the correct factor pair (p^*, q^*) with $p^* q^* = N$ and both factors prime, at $E_0 = 0$. Hence

$$\Delta E = E_1 - E_0 = 1 - 0 = 1 \quad (\text{arbitrary energy units}),$$

independent of N. The logarithmic term is subdominant in this convention: it vanishes on the ground state and adds non-negative energy elsewhere, so it cannot lower the gap below the value fixed by the integer-valued penalties.

4.2 Measurement time from the spectroscopic transition rate

From [2] the two-photon transition rate is $\Gamma = 70\ \mathrm{s}^{-1}$. The measurement time needed to resolve one transition event is

$$T_{\text{meas}} = \frac{1}{\Gamma} = \frac{1}{70\ \mathrm{s}^{-1}}.$$

Arithmetic: $1/70 = 0.0142857\ldots$, so

$$T_{\text{meas}} \approx 1.43 \times 10^{-2}\ \mathrm{s} \approx 14\ \mathrm{ms}.$$

A single spectroscopic interrogation at the demonstrated rate therefore distinguishes adjacent energy levels in roughly $14\ \mathrm{ms}$.

4.3 Experimental budget comparison

From [1] the analog design benchmark reports 50 tasks, 15 configurations per task, 3 rollouts per configuration, and a 2 h budget per attempt. The total number of attempts is

$$A = 50 \times 15 \times 3.$$

Arithmetic: $50 \times 15 = 750$; $750 \times 3 = 2{,}250$ attempts. The total wall-clock budget is

$$T_{\mathrm{budget}} = 2{,}250 \times 2\ \mathrm{h} = 4{,}500\ \mathrm{h}.$$

Converted to days: $4{,}500/24 = 187.5$ days. The ratio of budget to a single gap-resolving measurement is

$$\frac{T_{\mathrm{budget}}}{T_{\mathrm{meas}}} = \frac{4{,}500 \times 3{,}600\ \mathrm{s}}{1.4286 \times 10^{-2}\ \mathrm{s}} = \frac{1.62 \times 10^{7}\ \mathrm{s}}{1.4286 \times 10^{-2}\ \mathrm{s}} \approx 1.13 \times 10^{9}.$$

Arithmetic: $4{,}500 \times 3{,}600 = 1.62 \times 10^{7}$; $1.62 \times 10^{7} / 1.4286 \times 10^{-2} = 1.134 \times 10^{9}$. Resolving a constant gap is therefore negligible against the aggregate effort the analog design community already expends on design-space exploration.

4.4 Hidden-cost quantification under the smooth-encoding convention

We now drop the discrete product penalty and engineer only the logarithmic constraint $C(p,q;N) = (\ln p + \ln q - \ln N)^2$ with a single coefficient α, so that the gap is $\Delta E = \alpha\, C_{\min}$, where $C_{\min}$ is the minimal non-zero value of C over integer pairs.

4.4 Hidden-cost quantification under the smooth-encoding convention (continued)

The minimal non-zero value of C over integer pairs is attained by the pair with product closest to, but not equal to, N. Since p and q are integers, pq is an integer, and the closest integer to N other than N itself is $N \pm 1$. For that pair,

$$C_{\min} = \big(\ln(N \pm 1) - \ln N\big)^2 = \ln^2\!\Big(1 \pm \tfrac{1}{N}\Big).$$

For $N \gg 1$, $\ln(1 \pm 1/N) = \pm\frac{1}{N} - \frac{1}{2N^2} + O(N^{-3})$, so to leading order

$$C_{\min} \approx \frac{1}{N^2}.$$

The gap is therefore

$$\Delta E = \alpha\, C_{\min} \approx \frac{\alpha}{N^2}.$$

Energetic protection requires $\Delta E \gtrsim k_B T$. We take $T = 10\ \mathrm{mK} = 10^{-2}\ \mathrm{K}$ (the temperature scale at which trapped molecular ions are cooled, per the techniques reported in [2]) and $k_B = 1.380649 \times 10^{-23}\ \mathrm{J\,K^{-1}}$ (our assumption, Section 3). Then

$$k_B T = 1.380649 \times 10^{-23}\ \mathrm{J\,K^{-1}} \times 10^{-2}\ \mathrm{K} = 1.380649 \times 10^{-25}\ \mathrm{J}.$$

For a 128-bit integer, $N = 2^{128}$. Arithmetic: $2^{10} = 1024$, $2^{128} = 2^{8} \times (2^{10})^{12} = 256 \times 1024^{12}$. Since $1024^{12} = 10^{36}$ to three significant figures ($\log_{10} 1024 = 3.0103$, $12 \times 3.0103 = 36.1236$, $10^{0.1236} = 1.329$), we get $2^{128} \approx 256 \times 1.329 \times 10^{36} \times 10^{0} \approx 3.40 \times 10^{38}$. Then

$$N^2 \approx (3.40 \times 10^{38})^2 = 1.156 \times 10^{77}.$$

The required coefficient is

$$\alpha \approx k_B T \times N^2 = 1.380649 \times 10^{-25}\ \text{J} \times 1.156 \times 10^{77} \approx 1.60 \times 10^{52}\ \text{J}.$$

For comparison, [7] reports a Protection Deficit of order 10^{12} (dimensionless) for actively error-corrected quantum computation. The ratio is

$$\frac{1.60 \times 10^{52}\ \text{J}}{10^{12}} = 1.60 \times 10^{40},$$

i.e. the smooth-encoding coefficient exceeds the reported thermodynamic deficit scale by forty orders of magnitude. This is the quantitative statement that the exponential cost of analog factoring does not vanish under the smooth-encoding convention: it migrates from parameter precision to the physical magnitude of the engineered energy scale.

5. Results

The analysis yields four results, each derived with explicit arithmetic in Section 4.

R1 (Constant gap, discrete-penalty convention). With unit-coefficient integer-valued penalties, the sculpted factoring Hamiltonian has $E_0 = 0$, $E_1 = 1$, and $\Delta E = 1$ (arbitrary energy units), independent of N (Section 4.1). Correctness is then a thermal condition $k_B T \ll \Delta E$, not an analog-precision condition.

R2 (Measurement time). At the transition rate $\Gamma = 70\ \text{s}^{-1}$ reported in [2], one gap-resolving interrogation takes $T_{\text{meas}} = 1/\Gamma \approx 1.43 \times 10^{-2}$ s ≈ 14 ms (Section 4.2).

R3 (Budget comparison). The analog design benchmark of [1] aggregates $2{,}250$ attempts at 2 h each, i.e. $4{,}500$ h $= 1.62 \times 10^7$ s; the ratio of this budget to a single gap-resolving measurement is $\approx 1.13 \times 10^9$ (Section 4.3). Resolving a constant gap is negligible against the design-exploration effort the analog community already expends.

R4 (Hidden exponential cost, smooth-encoding convention). With only the logarithmic constraint engineered, $C_{\min} \approx 1/N^2$ and the gap is $\Delta E = \alpha/N^2$. Maintaining $k_B T$ protection at $T = 10$ mK for a 128-bit integer requires $\alpha \approx 1.60 \times 10^{52}$ J, forty orders of magnitude above the 10^{12}-scale Protection Deficit reported in [7] (Section 4.4).

Together, R1–R4 locate the exponential burden precisely: it is absent from the spectral gap under the discrete-penalty convention, and it reappears as the coefficient $\alpha \propto N^2$ under the smooth-encoding convention. The constant-gap claim is therefore conditional on the physical availability of integer-valued arithmetic penalties with N-independent precision.

6. Discussion

The central result — a constant gap $\Delta E = 1$ under the discrete-penalty convention — is a statement about an idealized Hamiltonian, and its practical significance depends entirely on whether the integer-valued penalties can be physically realized. We now argue against the result.

The implementation assumption is the whole result. Section 3 states the load-bearing assumption explicitly: the product penalty $|pq - N|^2$ must be implemented with unit coefficient and N-independent precision. Any physical implementation of $|pq - N|$ is itself a computation embedded in an analog substrate. If that substrate computes pq with analog components, the precision requirement has merely been relocated into the multiplier circuitry, and the exponential-precision objection returns one level down. The discrete-penalty convention is honest only if the arithmetic is digi-

tal at the point of energy comparison — for example, if the device natively represents integers as discrete charge or flux states whose product is exactly an integer. Whether such a native integer-arithmetic substrate exists at the required scale is an open question, and R1 must be read as conditional on it.

Falsification. The main result would be falsified by either of two observations. First, a demonstration that any physical realization of the penalty $|pq - N|^2$ requires coefficient precision or energy scale growing with N — this would collapse the discrete-penalty convention into the smooth-encoding convention, for which R4 already quantifies the cost as $\alpha \propto N^2$. Second, a measurement showing that the realized first excited state of a sculpted Hamiltonian sits below $E_0 + 1$ energy unit for large N, e.g. because the logarithmic term or cross-terms between penalties create spurious low-lying states that the case analysis of Section 4.1 does not cover. The case analysis is exhaustive over integer pairs, so the second failure mode can only arise from implementation error, not from the derivation itself.

Limitations. The result rests on the following assumptions, with sensitivity noted for each:

1. **Unit-coefficient integer penalties are physically available** (Section 3). The entire gap result R1 depends on this; if penalties are only smooth functions of continuous variables, R4 applies and the cost is $\alpha \approx 1.60 \times 10^{52}$ J for 128 bits.
2. **The spectroscopic rate transfers across systems.** We use $\Gamma = 70\ \mathrm{s}^{-1}$ from trapped molecular ions [2] as a proxy interrogation rate. If a sculpted device admits faster readout, R2 improves proportionally; if slower, it degrades. R2 is a timescale anchor, not a device claim.
3. **Thermal equilibrium and no tunneling** (Section 3). If the device relaxes non-thermally or tunnels between adjacent energy levels, the protection condition $k_B T \ll \Delta E$ is insufficient, and the gap alone does not guarantee correctness.
4. k_B **taken as a standard constant by our own assumption**, since no source in this analysis supplies it; the numerical value $1.380649 \times 10^{-23}\ \mathrm{J\,K^{-1}}$ is the defined SI value and its use is uncontroversial, but we flag it as an input not drawn from the cited works.
5. **The budget comparison mixes categories.** The $4{,}500$ h figure of [1] measures agentic design effort on transistor-level circuits, not the cost of building a sculpted Hamiltonian; the ratio 1.13×10^9 in R3 is a scale comparison, not a cost prediction.
6. **The smooth-encoding scaling is leading-order.** $C_{\min} \approx 1/N^2$ neglects the $O(N^{-3})$ correction in $\ln(1 \pm 1/N)$; for $N = 2^{128}$ the correction shifts α by a relative factor of order 10^{-38}, negligible here but relevant for small N.
7. **Primality penalties assume an oracle.** The term $(\mathrm{is_prime}(p) - 1)^2$ assumes primality is checkable within the device at zero additional gap cost. A physical primality test is itself a computation whose resource cost is not analyzed here.

Open questions. Three remain. (i) Does a native discrete integer-arithmetic substrate exist whose penalty terms are exactly integer-valued at scale? (ii) Does the fast-forwarding trade identified in [4] — where increasing the Hamiltonian norm shortens measurement time but the norm itself becomes the resource — admit a middle regime in which a moderately scaled α suffices for factoring integers of practical size (e.g. 64 bits rather than 128)? (iii) Do the non-Archimedean state spaces of [9], in which ultrametric topologies suppress linear error accumulation in simulation, offer a route to discrete penalty implementation that the Archimedean analog setting cannot? The simulation results of [9] report passive geometric fault tolerance sustained over 10^5 operational cycles with no continuous logical degradation, which is suggestive for the discrete-state question but was obtained in a Boolean threshold simulation, and its own text notes that physical substrates would experience sub-threshold leakage with fidelities of roughly 0.999.

Relation to prior conclusions. The topological-hydrodynamics analysis of [10] concludes that for cryptographic integers ($N > 10^{50}$) the energy gap required for analog resolution falls orders of magnitude below the thermal noise floor of any laboratory system at $T \geq 1$ mK. Our R4 is consistent with and quantifies that conclusion for the smooth-encod-

ing route: the required coefficient, not the gap itself, is what grows, and the growth is N^2. The two works agree that the smooth route is thermodynamically closed at cryptographic sizes; they differ in that we identify the discrete-penalty route as the remaining candidate and state the condition under which it evades the objection.

7. Practical Implications and Engineering Pathways

(a) Users and decisions. The result serves three groups. Researchers designing analog or Ising-type optimization hardware (the Class A systems of [8]) can use R1–R4 to decide at the design stage whether their penalty terms are discrete or smooth: if smooth, the coefficient budget grows as N^2 and the design is closed at cryptographic sizes; if discrete, the thermal condition $k_B T \ll \Delta E$ replaces the precision condition and the design remains open. Funding and roadmap decisions in quantum and analog computing can use R4 as a quantitative rejection criterion for smooth-encoding factoring proposals. Compiler and toolchain builders for physical computation can treat the discrete-versus-smooth distinction as a compilation check, analogous to type checking: the Hamiltonian is compilable to a protected device only if its penalties are integer-valued.

(b) Readiness level. This work is purely analytical and has not been peer reviewed; against the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and reported; concept and application formulated). No hardware component of the proposed sculpted Hamiltonian has been built, and no claim of manufacturing readiness is made.

(c) Smallest falsifying experiment. The minimal prototype is a two-penalty sculpted landscape with a small semiprime, e.g. $N = 35$, implemented in a platform with natively discrete integer states (superconducting flux qubits or trapped-ion vibrational modes). The experiment would measure the energy spectrum of the engineered Hamiltonian and check two numbers: that the ground state is the correct factor pair at $E_0 = 0$, and that the first excited state lies at $E_1 = 1$ energy unit. The outcome that refutes the discrete-penalty claim is a measured $E_1 < 1$ energy unit, or a measured dependence of $E_1 - E_0$ on N when the experiment is repeated for $N = 35, 77, 143$. A confirming outcome at small N does not prove the claim at cryptographic N; it establishes only that the penalty implementation is discrete at that scale, which is precisely the assumption whose scalability is open.

(d) Binding constraints. In the order in which they bind: first, the availability of a physical degree of freedom whose states are exactly integers and whose multiplication is exact (materials and device architecture — open question, no grounding quantifies it); second, the thermal budget, since protection requires $k_B T \ll \Delta E$ and the mK-scale cooling demonstrated for trapped ions [2] sets the reference temperature; third, readout rate, anchored at $\Gamma = 70\ \mathrm{s}^{-1}$ [2], which sets $T_{\mathrm{meas}} \approx 14$ ms per interrogation and bounds how many configurations can be interrogated per unit time; fourth, the Hamiltonian norm, since the fast-forwarding analysis of [4] shows that norm and measurement time trade against each other and that the norm is a real resource, and the sparse-simulation bound of [5] retains $\|H\|$ in its complexity $O\big(\tau \frac{\log(\tau/\varepsilon)}{\log\log(\tau/\varepsilon)}\big)$, so any realization or simulation with $\|H\| \propto N^2$ inherits that scaling.

(e) Order-of-magnitude projections. All inputs here are derived in Section 4 or cited. Projection: a smooth-encoding device protecting a 128-bit factoring instance at 10 mK requires an energy scale $\alpha \approx 1.60 \times 10^{52}$ J (Section 4.4); for comparison, the aggregate design-exploration budget of [1] is 1.62×10^7 s, and a single gap interrogation at the [2] rate is 1.43×10^{-2} s. These figures assume the smooth-encoding convention, $T = 10$ mK, $k_B = 1.380649 \times 10^{-23}\ \mathrm{J\,K^{-1}}$, and leading-order scaling $C_{\min} = 1/N^2$; uncertainty in $C_{\min}$ at the $O(N^{-3})$ level is negligible at this N. No cost figure in currency is derivable from the grounding and none is given.

(f) Risks and misuse paths. No direct misuse path is identified. The result, if it holds, does not enable factoring of deployed RSA keys: it states a condition under which an analog factoring device would be thermodynamically protected, and the open question of discrete penalty implementation at scale stands between the condition and any working de-

vice. Conversely, the result could be misused rhetorically by citing R1 while omitting the implementation assumption of Section 3; the correct reading is that R1 and R4 are a pair, and neither stands alone.

8. Conclusion

We analyzed whether constant-gap sculpted Hamiltonians evade the exponential-precision objection to analog computation. Under the discrete-penalty convention, the sculpted factoring Hamiltonian with unit-coefficient integer-valued penalties has ground-state energy $E_0 = 0$, first excited energy $E_1 = 1$, and a constant gap $\Delta E = 1$ independent of N; at the experimentally reported transition rate of $70\ \mathrm{s}^{-1}$ [2], one gap-resolving interrogation takes $14\ \mathrm{ms}$, negligible against the $4{,}500\ \mathrm{h}$ aggregate budget of the analog design benchmark [1]. Under the smooth-encoding convention, the minimal non-zero constraint value scales as $1/N^2$, and thermal protection at $10\ \mathrm{mK}$ for a 128-bit integer requires a coefficient $\alpha \approx 1.60 \times 10^{52}$ J, forty orders of magnitude above the 10^{12}-scale Protection Deficit reported for actively error-corrected quantum computation [7]. The exponential cost of analog factoring therefore does not vanish; it migrates from parameter precision to the physical implementation of arithmetic penalties. The constant-gap claim survives explicit analysis only under the discrete-penalty convention, and its viability reduces to a single testable question: whether integer-valued penalty terms with N-independent precision can be physically realized at scale. The smallest experiment that answers it — a spectral measurement of a sculpted two-penalty landscape at small semiprimes, repeated across increasing N — is specified in Section 7, and its outcome either establishes the first concrete instance of precision-free analog computation or closes the discrete-penalty route and, with it, the constant-gap program.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.45.4). 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: the execution sandbox was unavailable, so a second language model recomputed 5 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 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] Long-Horizon Analog Design Bench: Benchmarking Agents on Hours-Long Analog and Mixed-Signal Circuit Design Tasks. arXiv:2609.33356v1. https://arxiv.org/abs/2609.33356v1 [2] Precision Spectroscopy of Molecular Hydrogen Ions: Towards Frequency Metrology of Particle Masses. arXiv:physics/0612106v1.

https://arxiv.org/abs/physics/0612106v1 [3] Estimation of the directional parameter of the offset exponential and normal distributions in three-dimensional space using the sample mean. arXiv:1505.04647v1. https://arxiv.org/abs/1505.04647v1 [4] Fast-forwarding of Hamiltonians and Exponentially Precise Measurements. arXiv:1610.09619v2. https://arxiv.org/abs/1610.09619v2 [5] Exponential improvement in precision for simulating sparse Hamiltonians. arXiv:1312.1414v2. https://arxiv.org/abs/1312.1414v2 [6] Constructing Fermionic Hamiltonians with Non-Gaussianic low-energy states. arXiv:2502.15368v1. https://arxiv.org/abs/2502.15368v1 [7] https://papers.qnfo.org/papers/thermodynamic-viability-and-the-universality-of-feynman-matter. QNFO: Thermodynamic Viability and the Universality of Feynman Matter. [8] https://papers.qnfo.org/papers/programming-the-vacuum. QNFO: Programming the Vacuum. [9] https://papers.qnfo.org/papers/computational-simulation-approach-to-non-archimedean-quantum-architectures. QNFO: Computational Simulation Approach to Non-Archimedean Quantum Architectures. [10] https://papers.qnfo.org/papers/topological-hydrodynamics-and-the-spectral-gap. QNFO: TOPOLOGICAL HYDRODYNAMICS AND THE SPECTRAL GAP.