

Ultrametric Quantum Computation for Direct Integer Factorization: A Resource Analysis with a Largely Negative Verdict

Ultrametric Quantum Computation for Direct Integer Factorization: A Resource Analysis with a Largely Negative Verdict

Abstract

Factoring large integers underpins RSA, and the best-known polynomial-time quantum method, Shor's algorithm, extracts the period of modular exponentiation through a quantum phase-estimation detour that dominates circuit depth. A conjecture in the ultrametric research program holds that a quantum computer whose native state space is p -adic rather than Archimedean could access that period directly, bypassing phase estimation. This paper formalizes the conjecture as two concrete register models — a tensor-product adelic register and a Bruhat–Tits tree register — and costs every component exactly. The toy adelic register of dimension 810,000 admits a structured transform costing 2,888 operations against a dense baseline of 6.561×10^{11} , a structural separation of 2.27×10^8 ; we reproduce this arithmetic and retain the source's own caveat that it compares transform structures, not physics. For the tree register at a b -bit modulus with $p = 2$ and $d = 2b$, the register has $3 \cdot 2^{2b} - 2$ vertices and the transform costs $3(3 \cdot 2^{2b} - 2)$ operations, against $b(2b + 1)$ quantum Fourier transform gates in Shor's control register: the ratio is 4336.9 at $b = 8$ and approximately 10^{1227} at $b = 2048$. We further give a structural obstruction: p -adic valuations do not determine residues modulo the period, so no valuation-only measurement can replace phase estimation. The direct ultrametric route is therefore exponentially more expensive unless an unstated compression mechanism exists, which we state as the open conjecture.

1. Introduction

RSA and related cryptosystems derive their security from the assumed hardness of integer factorization, and every organization that deploys RSA today faces a concrete planning problem: if a quantum computer capable of factoring b -bit moduli in polynomial time is built, keys must be replaced on a schedule set by that machine's arrival. The canonical polynomial-time quantum method, Shor's algorithm, reduces factoring to period-finding — finding the multiplicative order r of a unit a modulo N — and executes it by preparing a superposition over exponents, computing $a^z \bmod N$ into a second register, and applying the quantum Fourier transform (QFT) over roughly N^2 points to estimate an eigenphase from which r is recovered by continued fractions. The phase-estimation detour is the entire quantum content of the algorithm: it converts a periodicity in the exponent into a measurable frequency. Any formulation of period-finding that removes this detour, or replaces it with something cheaper, would change both the gate-count projections used in quantum-threat assessments and the design targets of quantum hardware programs.

Present-day hardware cannot run Shor's algorithm at cryptographically relevant sizes regardless. The Tianyan cloud platform operates superconducting processors with 105 qubits at single-qubit, two-qubit, and readout fidelities of 99.90%, 99.56%, and 98.7% respectively, and demonstrates random circuit sampling on a 74-qubit region as a demonstration of quantum-advantage capability [2]; Devices of this class remain in the noisy intermediate-scale regime, and the depth of phase estimation is one of the dominant obstacles. A computational model that bypasses phase estimation would dramatically reduce circuit depth if its other costs did not grow faster.

The ultrametric research program motivates exactly such a replacement. The adelic-physics argument is that the choice of the real numbers $\mathbb{R}$ as the base field of physics is inherited from nineteenth-century analysis rather than from experiment, and that Ostrowski's theorem forces the p -adic completions $\mathbb{Q}_p$ to be treated as equally legitimate completions of the rationals [12]. The ultrametric-quantum-computation program proposes Bruhat–Tits trees as a physical substrate, Hecke operators as gates, and a threshold theorem for fault tolerance on the tree [10]. A companion register analysis computes exact gate counts for a finite adelic toy register and reports a structural separation of approximately 2.27×10^8 between a structured

transform and a dense baseline, while explicitly cautioning that this separation is a statement about transform structure, not about physics [11]. The conjecture at issue here is sharper: that an "adelic quantum computer" could find the period of $a^z \bmod N$ directly on the tree, without eigenvalue estimation.

This paper tests that conjecture by building the smallest complete models of it and costing every component. The outcome is largely negative, and the negative result is the useful one: it identifies precisely which step of the direct approach is exponentially costly, converting a sweeping ontological claim into a bounded technical statement.

Contributions.

1. Formal definitions of two ultrametric register models — the tensor-product adelic register and the Bruhat–Tits tree register — with exact dimension and vertex counts verified by shown arithmetic (Sections 3 and 4).
2. An exact resource comparison showing that direct p -adic period-finding on the tree register for a b -bit modulus costs $3(3 \cdot 2^{2b} - 2)$ transform operations against $b(2b + 1)$ QFT gates in Shor's control register, a ratio of 4336.9 at $b = 8$ and approximately 10^{1227} at $b = 2048$ (Sections 4 and 5).
3. A structural obstruction argument: the vector of p -adic valuations of a measured integer does not determine its residue modulo the period r , so no measurement reading only ultrametric magnitudes can extract r (Section 4, discussed in Section 6).
4. A full reproduction of the toy adelic register arithmetic — dimension 810,000, structured cost 2,888, dense baseline 6.561×10^{11} , separation 2.27×10^8 — with the dense baseline re-derived as a dense matrix-vector product of size dim^2 , and with the structural-versus-physical caveat made explicit (Sections 4 and 5).
5. Falsification criteria and a minimal de-risking pathway, combining a classical exact-cost calculation with a three-qubit p -adic phase-gate prototype (Sections 6 and 7).

2. Background and Related Work

The mathematical foundation for taking ultrametric state spaces seriously is well developed. [1] proves ultrametric analogues of three classical metric-geometry results: the Arens–Eells isometric embedding theorem, the Hausdorff extension theorem, and the Niemytzki–Tychonoff characterization of compactness, working over ultra-normed modules over arbitrary commutative rings and replacing the Dugundji extension theorem with a zero-dimensional Michael continuous selection theorem. For this paper, [1] establishes that the analytic infrastructure — embeddings, extensions, and function spaces over ultrametric spaces — exists and is rigorous, so defining a register on a Bruhat–Tits tree encounters no metric-geometric obstruction. What [1] does not provide is any notion of dynamics or unitary evolution; the obstacle examined here lies exactly there.

On the computational side, the direct-alternative-to-Shor genre has prior art. [3] constructs a quantum simulation of factorization by following Feynman's prescription: it defines a classical energy functional $E\{j; x_\sigma\} = \pi(x_\sigma)\pi(N_\sigma/x_\sigma)/j^2$ on a neighborhood $\mathcal{F}(j)$ of the input N , quantizes it with $p \rightarrow -i\partial_q$, and builds a computer state $\Psi(q; j)$ as a superposition over the dimension $F(j)$, drawing an explicit analogy to the Berry–Keating operator conjectured to realize the Hilbert–Pólya program for Riemann zeros. [3] thus shares this paper's target — factoring by a route other than QFT phase estimation — and also shares its epistemic status: the construction is exact in a stated range of validity but no asymptotic advantage over Shor is proven. Our work differs by costing the proposed alternative exactly and reporting where the cost explodes.

The universality question constrains what an ultrametric computer could even be. [8] proves that almost every two-qubit gate is universal, by showing that the non-universal gates form a lower-dimensional subset of the 16-dimensional manifold $\mathbf{U}(4)$, and states the complementary negative fact on which we rely: no gate that maps computation-basis states to computation-basis states of a fixed basis — "classical" gates in its terminology — can be universal, because such gates never create superposition. An ultrametric computer whose admissible unitaries are isometries of the tree, and therefore permute tree vertices, satisfies exactly this classicality condition with respect to the vertex basis. [8] also argues that the proper definition of universality depends on the laws of physics of the underlying technology, which is precisely the dependency this paper makes explicit for a p -adic technology.

Precedent for structure-protected computation that changes overhead but not power comes from measurement-based schemes. [7] considers qubits on which one can measure whether a pair is in the singlet or triplet state, protected by $SU(2)$ -invariance of the Hamiltonian, and proves (Theorem 1) that these protected measurements plus single-qubit X, Z unitaries approximate the full Clifford group, Pauli measurements, and T gates with polylogarithmic overhead in the error, via magic-state distillation. The lesson for the ultrametric program is that a protected or structured substrate re-derives the standard gate set at a quantified overhead; it does not yield a new computational model unless the substrate supports operations outside the standard power class.

Hardware context bounds what any new register model must eventually map onto. [2] describes the Tianyan cloud platform operating superconducting processors with 105 qubits at the fidelities quoted above, and demonstrates random circuit sampling on a 74-qubit region as a demonstration of quantum-advantage capability. Two features of [2] matter here: the connectivity is a fixed chip architecture adapted to the circuit layout rather than a tree [2]. Any ultrametric register must either be embedded in such hardware with an embedding overhead or motivate new fabrication.

The adelic viewpoint that motivates the conjecture has serious mathematical exemplars. [4] proves an arithmetic Hodge index theorem for adelic line bundles over projective varieties over finitely generated fields, extending results from number fields, and applies it to prove rigidity of preperiodic points of polarized algebraic dynamical systems, using nefness of admissible extensions and a fundamental inequality on Minkowski heights. [4] demonstrates that packaging local data at every place into a single adelic object yields global rigidity statements unreachable from any one completion — the exact pattern the ultrametric-computation conjecture hopes to exploit for computation. [6] supplies the measure-theoretic piece: it proves equidistribution of Galois-invariant sequences of points of small height on the Berkovich compactification of $\mathbb{P}^1$ with respect to quasi-adelic measures, showing that the μ_v -energy $\frac{1}{2|S_n|^2} \sum_{x \neq y} g_{\mu_v}(x, y)$ converges to zero at every place including the non-Archimedean ones. [6] thus shows that probability measures on p -adic analytic spaces are well behaved and admit limit theorems — a prerequisite for any measurement theory on a tree register — but the content is analytic number theory, not computational dynamics.

The proximate source of the conjecture is the QNFO program itself. [9] reports five research lines converging on an "Adelic Core" of valuation theory, Bruhat–Tits geometry, Ostrowski's theorem, and the adelic ring, and is explicit about its own epistemic status: the framework's physical claims are marked speculative, with no p -adic experimental signature confirmed, while remaining falsifiable through quantitative predictions. [10] develops the ultrametric quantum computation thesis in detail, states Ostrowski's theorem as the anchor — every nontrivial absolute value on $\mathbb{Q}$ is equivalent to the Archimedean one or to a p -adic one — and formulates a Threshold Theorem conjecture for fault tolerance on the tree, together with an explicit falsification list that includes the failure of hierarchical couplings to decay exponentially with depth. [11] is the closest methodological precedent: it builds a finite adelic toy register of dimension 810,000 for the prime set $S = \{2, 3, 5\}$ with local dimension $k_p = 4$, costs a structured functoriality transform at 2,888 operations against a dense baseline of 6.561×10^{11} , and states in its discussion that the separation factor 2.27×10^8 compares transform structures on the same register and says nothing about whether the required encodings exist or are efficiently preparable. We adopt [11]'s discipline: every number below is derived with shown arithmetic, and structural separations are kept strictly distinct from physical claims. Finally, [5] illustrates the reverse projection at the protocol level: its continuous-variable direct secure communication schemes encode discrete three-bit symbols by partitioning the real line into eight intervals and applying displacement operators $D(\beta)$ with $\beta = r + ir$, so that discrete information rides on a continuous amplitude read out by homodyne detection. This is the continuous screen carrying discrete structure — the mirror image of the conjecture's claim that quantum superposition is a discrete tree projected onto a continuous screen — and it shows that both directions of projection are implementable in principle; the open question is which direction carries the computational advantage.

3. Methods

3.1 Two register models

We formalize the conjecture with two models of increasing ambition, so that every cost quoted later attaches to a precisely defined object.

Model A: tensor-product adelic register. Fix a finite set of primes S and local exponents k_p . The global Hilbert space is

$$\mathcal{H}_S = \bigotimes_{p \in S} \mathbb{C}^{p^{k_p}},$$

with computational basis labelled by tuples of local indices. This is the register of [11]; we instantiate $S = \{2, 3, 5\}$ and $k_p = 4$ for all three primes, giving $\dim \mathcal{H}_S = 810,000$ (Section 4.1). Amplitudes are complex numbers; the adelic structure enters through the factorization of the transform into local p -adic Fourier transforms and cross-prime coupling stages.

Model B: Bruhat–Tits tree register. Fix a prime p and a depth d . The register is the vertex set of the Bruhat–Tits tree $T_{p,d}$ of $\mathrm{PGL}_2(\mathbb{Q}_p)$ truncated at depth d : a rooted tree whose root has $p + 1$ children and in which every vertex at depth $1 \leq \ell < d$ has p children. The computational basis is $\{|v\rangle : v \in V(T_{p,d})\}$ and the state space is $\mathbb{C}^{V(T_{p,d})}$ with the standard inner product. The vertex count is

$$V(p, d) = 1 + (p + 1) \sum_{\ell=0}^{d-1} p^\ell = 1 + (p + 1) \frac{p^d - 1}{p - 1},$$

which for $p = 2$ simplifies to $V(2, d) = 3 \cdot 2^d - 2$ (verified in Section 4.5).

3.2 The p -adic Fourier transform

For a function $f : S \rightarrow \mathbb{Q}_p$, the p -adic Fourier transform is

$$\tilde{f}(k) = \sum_{x \in S} f(x) \chi_p(kx),$$

where χ_p is the additive character of $\mathbb{Q}_p$. The conjecture's motivation is that $\tilde{f}$ is evaluated without eigenvalue estimation, because the character values are accessed directly as phase operations rather than through a phase-estimation subroutine. On Model A the transform factorizes into local transforms on each factor $\mathbb{C}^{p^{k_p}}$ plus cross-prime coupling stages; on Model B it is a hierarchical transform over the tree, applied at every vertex.

3.3 Period-finding without phase estimation: the claimed direct route

Given N and a coprime integer a , the modular exponentiation function $f(z) = a^z \bmod N$ has period r . In the direct ultrametric route one prepares the uniform superposition $\frac{1}{\sqrt{|S|}} \sum_{z \in S} |z\rangle$, applies the oracle $U_f : |z\rangle|0\rangle \mapsto |z\rangle|f(z)\rangle$, and then performs the p -adic transform on the first register, with the claim that measurement yields enough information to recover r classically. The critical question — which the resource analysis answers — is what the register must store and what the measurement must read for this to work. We assume, as the direct analogue of Shor's control register of $2b$ qubits spanning $\{0, \dots, N^2 - 1\}$, that the tree register is given depth $d = 2b$ so that its leaf count matches the scale of the classical control register; the choice $d = 2b$ is the most charitable one for the conjecture, since any smaller depth cannot resolve the period by the same counting argument that fixes Shor's register size.

3.4 Costing conventions

Three conventions fix every number below. First, an elementary operation is one two-level unitary or one arithmetic step in a transform butterfly. Second, a structured transform on a tensor factor of dimension $q = p^{k_p}$ implemented by a radix- p Cooley–Tukey algorithm costs $q \log_p q = k_p q$ operations. Third, a dense baseline on a register of dimension D is a dense matrix-vector product costing D^2 operations; Section 4.7 cross-checks this against a fast-transform baseline to show that the baseline convention, not the structured cost, sets the separation factor in Model A.

4. Analysis

4.1 Dimension of the toy adelic register

From [11] the register has $S = \{2, 3, 5\}$ and $k_p = 4$. Compute each local dimension:

$$2^4 = 16, \quad 3^4 = 81, \quad 5^4 = 625.$$

Multiply stepwise:

$$16 \times 81 = 1,296, \quad 1,296 \times 625 = 810,000.$$

Thus $\dim \mathcal{H}_S = 810,000$, matching the value reported in [11].

4.2 Structured transform cost on Model A

The local p -adic Fourier transform on a factor of size p^{k_p} costs $k_p p^{k_p}$ operations by the convention of Section 3.4. Summing over the three primes:

$$\begin{aligned} C_{\text{loc}} &= \sum_{p \in S} k_p p^{k_p} \\ &= 4 \cdot 2^4 + 4 \cdot 3^4 + 4 \cdot 5^4 \\ &= 64 + 324 + 2,500 \\ &= 2,888. \end{aligned}$$

This reproduces the structured cost $C_{\text{str}} = 2,888$ reported in [11]. In [11]'s own stage decomposition the total is split as $C_{\text{loc}} + C_{\text{couple}} + C_{\text{agg}} = 2,000 + 800 + 88$; the exact stage split is a bookkeeping convention, and the invariant quantity is the total 2,888, which the radix- p count above derives independently.

4.3 Dense baseline on Model A

A dense implementation treats the entire register as a flat vector of length $D = 810,000$ and applies an arbitrary unitary as a dense matrix-vector product:

$$C_{\text{dense}} = D^2 = 810,000^2.$$

Compute: $810,000 = 81 \times 10^4$, so

$$810,000^2 = 81^2 \times 10^8 = 6,561 \times 10^8 = 6.561 \times 10^{11}.$$

This reproduces the dense baseline 6.561×10^{11} reported in [11], and identifies its origin: it is exactly $\dim^2$, the cost of an unstructured unitary on the full register.

4.4 Separation factor on Model A

$$\frac{C_{\text{dense}}}{C_{\text{str}}} = \frac{6.561 \times 10^{11}}{2,888}.$$

Compute: $6.561 \times 10^{11} / 2,888 = 6.561 / 2,888 \times 10^{11} = 0.0022725 \times 10^{11} \approx 2.2725 \times 10^8$.

The separation factor is therefore $\approx 2.27 \times 10^8$, matching [11]. Following [11]'s own discussion, this number compares transform structures on the same register; it is not evidence that the required encodings exist or are efficiently preparable, and it is not a comparison against Shor's algorithm.

4.5 Vertex count of the tree register

The count $V(p, d)$ follows by summing vertices level by level. The root is a single vertex. The root has $p + 1$ children, so level 1 has $p + 1$ vertices. Every vertex at level $1 \leq \ell < d$ has p children, so level $\ell + 1$ has p times as many vertices as level ℓ : level 2 has $(p + 1)p$, level 3 has $(p + 1)p^2$, and in general level ℓ has $(p + 1)p^{\ell-1}$ vertices for $1 \leq \ell \leq d$. Summing a geometric series,

$$V(p, d) = 1 + (p + 1) \sum_{\ell=1}^d p^{\ell-1} = 1 + (p + 1) \frac{p^d - 1}{p - 1}.$$

For $p = 2$ the denominator is 1 and the sum is $2^d - 1$, so

$$V(2, d) = 1 + 3(2^d - 1) = 3 \cdot 2^d - 2.$$

As a check at small depth: $V(2, 1) = 3 \cdot 2 - 2 = 4$, which matches one root plus $p + 1 = 3$ children; $V(2, 2) = 12 - 2 = 10$, matching $1 + 3 + 6$. Both check out. The register analysis of [11] reports the vertex counts $V(2, 10) = 3070$ and $V(3, 5) = 485$; the formula reproduces both: $V(2, 10) = 3 \cdot 1024 - 2 = 3072 - 2 = 3070$, and $V(3, 5) = 1 + 4 \frac{3^5 - 1}{2} = 1 + 4 \cdot 121 = 485$.

4.6 Transform cost on the tree register and the ratio against Shor's control register

The hierarchical p -adic transform on Model B is applied at every vertex, and at each vertex it combines the amplitudes of the $p + 1$ branches by the butterfly of Section 3.2, which costs $p + 1$ operations per vertex for the character phases and sums. For $p = 2$ this is 3 operations per vertex, giving the total transform cost

$$C_{\text{tree}} = 3 V(2, d) = 3 (3 \cdot 2^d - 2).$$

With the charitable depth choice $d = 2b$ of Section 3.3,

$$C_{\text{tree}}(b) = 3 (3 \cdot 2^{2b} - 2).$$

On the Shor side, the control register holds $n = 2b$ qubits. We adopt the convention that the QFT on n qubits costs $\frac{n(n+1)}{2}$ two-qubit gates: this counts n Hadamard gates plus $\frac{n(n-1)}{2}$ controlled-phase rotations, with the final swap layer of $\lfloor n/2 \rfloor$ gates omitted, so $\frac{n(n+1)}{2}$ is the convention used throughout; alternative QFT decompositions change C_{Shor} by at most a constant factor (see the Limitations list in Section 6). Hence

$$C_{\text{Shor}}(b) = \frac{2b(2b+1)}{2} = b(2b+1).$$

Ratio at $b = 8$. Compute each side exactly:

$$2^{16} = 65,536, \quad 3 \cdot 65,536 = 196,608, \quad V(2, 16) = 196,606,$$

$$C_{\text{tree}}(8) = 3 \times 196,606 = 589,818,$$

$$C_{\text{Shor}}(8) = 8 \times 17 = 136.$$

The ratio is

$$\frac{589,818}{136} = 4336 + \frac{122}{136} \approx 4336.9.$$

Already at an 8-bit modulus — far below any cryptographic size — the direct tree route costs a factor of $\approx 4.3 \times 10^3$ more transform operations than the phase-estimation detour it is meant to replace.

Ratio at $b = 2048$ (RSA-2048 scale). Work in base-10 logarithms. The tree cost is dominated by $9 \cdot 2^{4096}$:

$$\log_{10} C_{\text{tree}} \approx \log_{10} 9 + 4096 \log_{10} 2 = 0.95424 + 4096 \times 0.30103 = 0.95424 + 1233.019 = 1233.973.$$

The Shor cost is

$$C_{\text{Shor}}(2048) = 2048 \times 4097 = 8,390,656, \quad \log_{10} C_{\text{Shor}} = 6.9238.$$

The logarithm of the ratio is

$$1233.973 - 6.9238 = 1227.049,$$

so

$$\frac{C_{\text{tree}}(2048)}{C_{\text{Shor}}(2048)} \approx 1.12 \times 10^{1227}.$$

The direct ultrametric route is therefore more expensive by a factor of order 10^{1227} at the modulus size relevant to RSA-2048. The exponent 1227 is itself the headline: the deficit grows as $\log_{10} C_{\text{tree}} - \log_{10} C_{\text{Shor}} \approx 2b \log_{10} 2 + \log_{10} 9 - \log_{10} 2 - 2 \log_{10} b \approx 0.602b + 0.653 - 2 \log_{10} b$ decades, so the gap widens essentially linearly in b , with coefficient 0.602 decades per bit, consistent with the $b = 2048$ computation, and without bound as moduli grow.

4.7 Cross-check: the dense baseline, not the structured cost, sets the Model A separation

Section 4.4 derived the separation 2.27×10^8 from a dense baseline D^2 . To confirm that this baseline convention, rather than an unusually small structured cost, drives the number, compare both against a fast full-register transform. A fast Fourier-type transform on a register of dimension D costs $D \log_2 D$. With $D = 810,000$: since $2^{19} = 524,288$ and $2^{20} = 1,048,576$, we have $\log_2 D \approx 19 + \log_2(810000/524288) \approx 19 + 0.626 = 19.63$, so

$$C_{\text{fast}} \approx 810,000 \times 19.63 \approx 1.59 \times 10^7.$$

Three ratios follow:

$$\begin{aligned} \frac{C_{\text{dense}}}{C_{\text{fast}}} &= \frac{6.561 \times 10^{11}}{1.59 \times 10^7} \approx 4.13 \times 10^4, & \frac{C_{\text{fast}}}{C_{\text{str}}} &= \frac{1.59 \times 10^7}{2,888} \approx 5.5 \times 10^3, \\ \frac{C_{\text{dense}}}{C_{\text{str}}} &= 4.13 \times 10^4 \times 5.5 \times 10^3 \approx 2.27 \times 10^8, \end{aligned}$$

which is internally consistent with Section 4.4. The structured cost 2,888 is genuinely small — it beats even the fast full-register transform by three and a half orders of magnitude — but the headline separation of 2.27×10^8 is set by charging the baseline as a dense D^2 product. This is exactly the caveat [11] states: the separation compares transform structures on the same register and says nothing about whether the required encodings exist or are efficiently preparable.

4.8 Structural obstruction: valuations do not determine residues modulo the period

The direct route of Section 3.3 requires the final measurement to yield enough information to recover r . A measurement that reads only ultrametric magnitudes reads the vector of p -adic valuations of the measured integer at the primes of the register. We now show this vector cannot determine the residue class modulo r , by explicit counterexample.

Take $N = 7$, $a = 5$. The powers of 5 modulo 7 are 5, 4, 6, 2, 3, 1, so the period is $r = 6$. Let the register's prime set be $S = \{2, 3\}$, the primes dividing r . Consider the two exponents $z_1 = 1$ and $z_2 = 5$. Their valuation vectors are

$$(v_2(1), v_3(1)) = (0, 0), \quad (v_2(5), v_3(5)) = (0, 0).$$

The vectors are identical: both 1 and 5 are units at 2 and at 3. But their residues modulo $r = 6$ differ, since $1 \not\equiv 5 \pmod{6}$. Consequently the map

$$z \mapsto (v_p(z))_{p \in S}$$

has fibers that contain multiple residue classes modulo r , and any measurement whose outcome is the valuation vector — the only information an ultrametric magnitude reading exposes — cannot distinguish z_1 from z_2 . Since recovering r from the period-finding measurement requires distinguishing residue classes modulo r (this is what continued-fraction post-processing consumes in Shor's algorithm), no valuation-only measurement can replace phase estimation. The obstruction is structural, not a matter of gate counts: it survives any improvement in the transform's efficiency, and it applies to Model A and Model B alike, because both models locate the readable information in ultrametric magnitudes. The direct route would need a measurement that reads p -adic *phase* information — characters $\chi_p(kx)$ — not merely p -adic *magnitude* information, and no such measurement is described in the ultrametric-quantum-computation program [10].

5. Results

Every number below is derived in Section 4 with shown arithmetic; none is measured or simulated.

Quantity	Value	Derivation
Toy adelic register dimension, $S = \{2, 3, 5\}, k_p = 4$	810,000	Section 4.1
Structured transform cost C_{str}	2,888	Section 4.2
Dense baseline $C_{\text{dense}} = D^2$	6.561×10^{11}	Section 4.3
Separation factor $C_{\text{dense}}/C_{\text{str}}$	$\approx 2.27 \times 10^8$	Section 4.4
Fast full-register baseline $D \log_2 D$	$\approx 1.59 \times 10^7$	Section 4.7
Tree vertex count $V(2, d)$	$3 \cdot 2^d - 2$	Section 4.5
Tree transform cost at b -bit modulus, $d = 2b$	$3(3 \cdot 2^{2b} - 2)$	Section 4.6
Shor control-register QFT cost at b -bit modulus	$b(2b + 1)$	Section 4.6
Cost ratio at $b = 8$	≈ 4336.9	Section 4.6
Cost ratio at $b = 2048$	$\approx 1.12 \times 10^{1227}$	Section 4.6
Valuation-obstruction counterexample	$N = 7, a = 5, r = 6; z = 1$ and $z = 5$ share valuation vector $(0, 0)$ at $S = \{2, 3\}$ but differ mod 6	Section 4.8

Three results carry the verdict. First, the toy register arithmetic of [11] reproduces exactly under independent derivation, confirming that the Model A numbers are theorems about a defined object. Second, the tree register comparison is exponentially unfavorable: the ratio $C_{\text{tree}}/C_{\text{Shor}}$ grows from 4336.9 at $b = 8$ to $\approx 10^{1227}$ at $b = 2048$, so the direct route is not merely worse at cryptographic scale but worse by an exponent that itself scales linearly in b . Third, the valuation obstruction of Section 4.8 shows that no measurement reading only ultrametric magnitudes can extract the period, independently of all gate counts. The two results are complementary: even if a measurement existed that evaded the obstruction, its transform would cost exponentially more than the phase-estimation detour it replaces.

6. Discussion

The central claim of this paper is deliberately narrow: the register dimensions, vertex counts, and gate counts above are theorems about two precisely defined models, while the conjecture that motivated them — that an adelic quantum computer factors by direct period access — remains untested at the level of physics. The negative verdict attaches to the models as defined, not to every model a p -adic technology could conceivably support.

What the result does not rule out. The costing assumes the tree transform must be applied at every vertex and costs $p + 1$ operations per vertex. A compression mechanism that evaluates the hierarchical transform on a subtree of size polylogarithmic in 2^{2b} , or that reads period information from a constant number of vertices, would change the ratio

$C_{\text{tree}}/C_{\text{Shor}}$ qualitatively. No such mechanism is described in [10] or [11]; stating its existence precisely is the open conjecture this paper isolates. Similarly, the valuation obstruction of Section 4.8 binds only measurements that read magnitudes; a measurement reading p -adic characters directly would evade it, and the burden of specifying one lies with the conjecture's proponents.

Falsification of the main result. The main result — exponential unfavorable costing of the direct route — would be falsified by either of two demonstrations. First, a constructive one: an explicit algorithm on Model B that recovers r for the pair $(N, a) = (7, 5)$ from a measurement whose outcome space has size polynomial in b , together with a transform cost on Model B asymptotically below $b(2b + 1)$; the toy case $(7, 5)$ is fully classically simulable, so this demonstration is checkable by exhaustive enumeration. Second, an analytical one: a proof that the hierarchical p -adic transform on $T_{2,2b}$ admits a sublinear-in- V evaluation preserving the character information needed for period recovery. If either is produced, the ratio $\approx 10^{1227}$ at $b = 2048$ no longer applies, and the direct route reopens.

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

1. *Depth choice* $d = 2b$. This is the most charitable choice; a larger d only increases C_{tree} , and a smaller d cannot resolve the period by the counting argument of Section 3.3. The conclusion is insensitive to the exact constant.
2. *Per-vertex transform cost* $p + 1$. If the per-vertex cost were $O(1)$ instead, C_{tree} would still scale as $V(2, 2b) \sim 3 \cdot 2^{2b}$, and the ratio would remain exponential. The conclusion is insensitive to the per-vertex constant.
3. *Shor-side cost* $b(2b + 1)$. This counts QFT gates only, under the $\frac{n(n+1)}{2}$ convention of Section 4.6, and excludes modular exponentiation, which dominates Shor's total cost; including it would only widen the gap. Alternative QFT decompositions change C_{Shor} by at most a constant factor, which does not affect the exponential verdict. The conclusion is conservative in Shor's favor.
4. *Operation convention*. Counting two-level unitaries and butterflies is standard; a convention charging tree-edge communication separately would increase C_{tree} further.
5. *Measurement model*. The obstruction of Section 4.8 assumes the measurement reads valuation vectors at the register's primes. A character-reading measurement is not modeled here and is not excluded by this paper's arithmetic; it is excluded only by the absence of any proposed mechanism in [10].
6. *No error correction*. All costs are noiseless. Fault-tolerance overheads on the tree are themselves conjectural ([10] states the Threshold Theorem as a conjecture with an argument sketch), so no corrected-cost comparison is possible on either side.

Open questions. Whether the Bruhat–Tits tree admits a fault-tolerant gate set at all is open — [10]'s own falsification list names the failure of hierarchical couplings to decay exponentially with depth as a killing negative result. Whether the adelic product formula, which [11] verifies numerically ($\prod_v |6|_v = 1$), can be exploited as a computational primitive rather than an identity is open. Whether the equidistribution theory of [6] on Berkovich spaces supports any measurement theory with nonzero information extraction about residues is open.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are the proponents and reviewers of the ultrametric quantum computation program itself: the analysis converts their conjecture into a bounded technical target — the missing compression mechanism and the missing character-reading measurement — and gives them two exact checkpoints (the $(7, 5)$ toy case and the $b = 8$ ratio of 4336.9) against which any proposed mechanism can be tested. A second audience is the quantum-threat assessment community: if the conjecture had held, gate-count projections for factoring would have needed a new entry; this paper's result says the existing phase-estimation-based projections remain the relevant ones, so no change to key-replacement planning is warranted by ultrametric considerations. Beyond these, the work is foundational: it defines and costs register models that did not previously have exact costs, and no deployment decision depends on it today.

(b) Readiness level. Against the NASA Technology Readiness Level definitions, this work is analytical research that has not been peer reviewed: it is TRL 1–2. Basic principles are formulated and the models are defined and costed on paper; no proof-of-concept hardware, simulation, or experimental validation exists. Analytical work of this kind cannot exceed TRL 3, and

nothing here approaches it: no component of either register model has been simulated, let alone built.

(c) Smallest next test. The smallest de-risking experiment is classical and costs only compute: exhaustively enumerate the period-finding circuit of Section 3.3 on the toy instance $N = 7$, $a = 5$, $r = 6$ over the register $S = \{2, 3\}$, and search over all measurements whose outcome space has size polynomial in b for one that recovers $r = 6$ from the post-transform state. The refuting outcome is a demonstrated polynomial-size measurement that outputs r correctly for all coprime a modulo $N = 7$; the confirming outcome is a proof (or exhaustive verification) that no such measurement exists on this register, which would upgrade the Section 4.8 obstruction from a counterexample to a classification. A minimal hardware companion, at the level of a tabletop prototype, is a three-qubit p -adic phase-gate demonstration: implement the character phases $\chi_p(kx)$ for $p = 2$, $k \in \{0, 1, 2\}$ on three qubits and measure whether the transform output matches the classical $\tilde{f}$ values; a mismatch at this scale would falsify the gate-level claims of [10] before any larger investment.

(d) Binding constraints. In order of first binding: (i) the measurement mechanism — no character-reading measurement exists even in proposal form, and this is the constraint that binds first and hardest; (ii) the encoding — [11] states that nothing is known about whether the required register encodings are efficiently preparable, and Section 4.4's caveat carries over verbatim; (iii) fabrication — any tree register must either embed in two-dimensional chip architectures of the kind [2] describes, with an unquantified embedding overhead, or motivate new fabrication for which no materials pathway is proposed; (iv) fault tolerance — the threshold theorem for the tree is a conjecture with an argument sketch, so error-correction overheads are unknown on both sides of the comparison. Data, compute, and cost drivers cannot be quantified from this paper's derivations and remain open questions.

(e) Order-of-magnitude bounds. The only quantitative bounds available are the ones derived here, and they are exact for the models as defined: the direct route costs $3(3 \cdot 2^{2b} - 2)$ transform operations against $b(2b + 1)$ for the Shor control register, a deficit of 4336.9 at $b = 8$ and $\approx 1.12 \times 10^{1227}$ at $b = 2048$. No cost, market, or performance figure beyond these model-internal numbers is derivable from this paper, and none is projected.

(f) Risks and misuse paths. There is no direct misuse path: the paper builds no cryptanalytic capability and breaks no scheme. The residual risk is rhetorical rather than technical — the structural separation 2.27×10^8 of Model A could be quoted without the Section 4.4 and 4.7 caveats as evidence that adelic registers outperform conventional ones, when in fact it compares a radix recursion against a dense baseline on the same register and is silent on encodings, preparation, and physics. This paper's own framing is designed to make that quotation impossible without distortion.

8. Conclusion

This paper tested the conjecture that an ultrametric quantum computer could factor integers by accessing the period of $a^z \bmod N$ directly, without the phase-estimation detour of Shor's algorithm. The test was conducted by building the smallest complete models of the conjecture and costing every component with shown arithmetic. The verdict is largely negative and precisely localized. On the toy adelic register, the reported arithmetic of [11] — dimension 810,000, structured cost 2,888, dense baseline 6.561×10^{11} , separation 2.27×10^8 — reproduces exactly, but the separation is a statement about transform structure, not physics. On the tree register at a b -bit modulus with the charitable depth $d = 2b$, the direct route costs $3(3 \cdot 2^{2b} - 2)$ transform operations against $b(2b + 1)$ QFT gates in Shor's control register: worse by 4336.9 at $b = 8$ and by $\approx 1.12 \times 10^{1227}$ at $b = 2048$. Independently of gate counts, a structural obstruction holds: p -adic valuations do not determine residues modulo the period, exhibited concretely on $(N, a) = (7, 5)$ with $r = 6$, so no measurement reading only ultrametric magnitudes can replace phase estimation. The direct ultrametric route is therefore exponentially more expensive as defined, and the conjecture survives only as a sharp, falsifiable open question: does a compression mechanism or a character-reading measurement exist that evades both the exponential costing and the valuation obstruction? The toy instance $(7, 5)$ is classically simulable and constitutes the minimal test. Until such a mechanism is exhibited, phase estimation remains the cheapest known quantum route to the period, and the gate-count projections used in quantum-threat assessments stand unchanged.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.6-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, zai-org/glm-5.3, openai/gpt-oss-120b, 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 6 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] An embedding, an extension, and an interpolation of ultrametrics. arXiv:2008.10209v2. <https://arxiv.org/abs/2008.10209v2> [2] Tianyan: Cloud services with quantum advantage. arXiv:2512.10504v2. <https://arxiv.org/abs/2512.10504v2> [3] Simulating Factorization with a Quantum Computer. arXiv:1505.00642v4. <https://arxiv.org/abs/1505.00642v4> [4] The arithmetic Hodge index theorem for adelic line bundles II. arXiv:1304.3539v2. <https://arxiv.org/abs/1304.3539v2> [5] Continuous variable direct secure quantum communication using Gaussian states. arXiv:1909.09697v2. <https://arxiv.org/abs/1909.09697v2> [6] Quasi-adelic measures and equidistribution on $\mathbb{P}^1$. arXiv:1502.04660v3. <https://arxiv.org/abs/1502.04660v3> [7] Symmetry Protected Quantum Computation. arXiv:2105.04649v3. <https://arxiv.org/abs/2105.04649v3> [8] Universality in Quantum Computation. arXiv:quant-ph/9505018v1. <https://arxiv.org/abs/quant-ph/9505018v1> [9] QNFO: Five Pillars, One Structure: Consilient Convergence in QNFO Research [10] QNFO: Ultrametric Quantum Computation and the Langlands Program [11] QNFO: Functoriality as a Transform: Register and Gate-Cost Analysis of a Langlands-Structured Adelic Quantum Register [12] QNFO: The Adelic Physics Program: Epistemological Foundations and Communications Framework