V1.0.0

Ultrametric Tree Connectivity as a Passive Fault-Tolerance Mechanism for Quantum Processors

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Fault-tolerant quantum computation currently relies on active stabilizer codes that demand repeated syndrome extraction, low-latency classical decoding, and large ancilla overhead; we assume a surface-code overhead in the range 10^2–10^4 physical qubits per logical qubit for comparison, while [9] estimates a $10^3{:}1$ ratio under its own assumptions. A passive protection mechanism that exploits hardware geometry rather than active decoding cycles would reduce this overhead, but no established design principle converts connectivity geometry directly into an error barrier. We investigate the hypothesis that arranging qubits on a Bruhat–Tits (ultrametric) tree confines correlated errors within subtrees through the strong triangle inequality. Using a stochastic depolarizing model with per-edge error probability ε, we derive an exponential barrier $B(d) = 2^d$ at tree depth d, the logical error bound $p_L(d) = \varepsilon^{B(d)}$, and the suppression factor $S(d) = \varepsilon / p_L(d)$. For $\varepsilon = 0.4$ we compute $p_L(3) = 6.5536 \times 10^{-4}$ and $p_L(4) = 4.2950 \times 10^{-7}$, giving suppression factors of 6.10×10^2 and 9.31×10^5. Term-by-term evaluation of a recursive 8-leaf majority vote yields failure probabilities 0.1736704, 5.55×10^{-3}, and 2.905×10^{-10} at successive levels, matching the values reported by prior barrier analyses. Ball-code size formulas give 5, 17, and 53 qudits for depths 1–3 on a ternary tree. At $\varepsilon = 0.4$, the model yields suppression factors of 6.10×10^2 at depth 3 and 9.31×10^5 at depth 4, and it defines a five-atom prototype experiment that would confirm or falsify the barrier model.

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.). Ultrametric Tree Connectivity as a Passive Fault-Tolerance Mechanism for Quantum Processors. QNFO Preprint v1.0.0.

---

1. Introduction We consider quantum processors whose logical-qubit overhead under active stabilizer codes we assume to lie in the range 10^2–10^4 physical qubits per logical qubit. The engineering cost of real-time syndrome extraction, low-latency classical decoders, and ancilla refresh cycles limits the size of computations that near-term devices can run, and it inflates power consumption and system footprint. Hardware architects on every major platform must therefore allocate a large fraction of their qubit budget to error correction machinery rather than to useful computation. If logical information could instead be stored in a geometry that *passively* blocks error propagation, the need for active correction cycles would shrink, freeing resources for deeper circuits and larger problem instances.

We consider the concrete scenario of a quantum computing platform in which qubits (or higher-dimensional qudits) are arranged on the vertices of a regular odd-p Bruhat–Tits tree. In such a layout, the strong triangle inequality of the ultrametric distance forces any error chain connecting two distant vertices to cross a cut whose size grows exponentially with the hierarchical level. Correlated error strings consequently become trapped inside subtrees, and logical information encoded at the root experiences an effective error barrier.

The decisive question is whether this geometric confinement can raise the effective fault-tolerance threshold of a concatenated code—here the $[[7, 1, 3]]$ Steane code—by a material factor relative to a planar nearest-neighbour layout at identical gate fidelities. If the answer is affirmative, quantum hardware designers could adopt tree-like interconnects to reach compa-

rable logical error rates with far fewer physical qubits, reducing overhead from the 10^2–10^4 range toward tens of qudits per logical qubit at modest depths.

Contributions.

1. We derive the exponential barrier scaling $B(d) = 2^d$ from the ultrametric strong triangle inequality and obtain the logical error bound $p_L(d) = \varepsilon^{B(d)}$, evaluated in closed form for depths $d = 3, 4$ (Section 4.1).
2. We compute suppression factors $S(3) = 6.10 \times 10^2$ and $S(4) = 9.31 \times 10^5$ with full arithmetic (Section 4.2).
3. We evaluate recursive 8-leaf majority-vote failure probabilities $p^{(1)} = 0.1736704$, $p^{(2)} \approx 5.55 \times 10^{-3}$, and $p^{(3)} \approx 2.905 \times 10^{-10}$ by exact binomial sums, reproducing the values reported in the prior barrier analysis [8] (Section 4.3).
4. We present closed-form ball-code sizes $n_t = 1 + (p+1)\frac{p^t-1}{p-1}$ and evaluate them for $p = 3$ at depths $t = 1, 2, 3$, giving 5, 17, and 53 qudits (Section 4.4).
5. We derive a thermal activation bound $p_{\text{thermal}}(3) = e^{-8} \approx 3.35 \times 10^{-4}$ under a unit-temperature Boltzmann model (Section 4.5).
6. We translate these analytical results into design guidelines, a technology-readiness assessment, and a minimal five-atom falsification experiment (Section 7).

The analysis also bears on adjacent problems in which hierarchical geometry governs information flow, including holographic quantum error-correcting codes, spin-glass ultrametricity, and p-adic models of condensed-matter systems.

2. Background and Related Work Continuous-time quantum error correction (CTQEC) treats noise and correction as simultaneous continuous processes and isolates a correctable subsystem $\mathcal{H}^A$ via a recovery CPTP map $\mathcal{R}$; the formalism is set out in [1]. Although CTQEC offers a unified dynamical picture of correction, it remains an *active* scheme and does not provide a passive geometric barrier to error propagation, which is the mechanism this paper studies.

Entanglement-assisted quantum error-correcting codes (EAQECCs) extend the stabilizer formalism by allowing pre-shared ebit pairs between encoder and decoder, thereby improving code rate or distance beyond what unassisted stabilizer codes achieve [2]. EAQECCs nonetheless still require active syndrome measurement and decoding, so they do not address the passive-protection goal that motivates the present work.

Classical error-correction theory supplies the majority-vote mechanism through the three-bit repetition code, in which a simple threefold redundancy corrects any single bit flip by majority decision [4]. This principle underlies the recursive majority-vote decoder we analyze on the tree, generalized from three votes to a wider leaf set.

The stabilizer formalism introduced by Shor and Steane provides distance-d codes such as the $[[7, 1, 3]]$ Steane code, in which any error can be expanded in a Pauli basis and corrected by syndrome extraction [5]. This formalism justifies representing each leaf error on the tree by a Pauli flip and applying majority-vote decoding over a branching structure; the concatenated scheme we analyze takes the Steane code as its base.

General quantum error-correction reviews establish that protecting quantum information requires encoding it into a subspace of a larger Hilbert space and extracting syndromes to identify and reverse errors [6]. Our approach replaces the syndrome-extraction step with a geometric barrier, so these reviews define the baseline architecture whose overhead we seek to reduce.

Ultrametric spaces satisfy the strong triangle inequality $d(x, y) \leq \max\{d(x, z), d(z, y)\}$, as formalized in the mathematical literature on ultrametric embeddings and extensions [7]. This inequality implies that any path connecting two distant vertices must cross a "cut" whose size grows with the hierarchical level; we exploit exactly this property to bound error-chain propagation in Section 4.1.

The QNFO barrier-analysis work on ultrametric passive error correction reports concrete numerical targets for a ternary tree: barrier heights $B(3) = 8$ and $B(4) = 16$; logical error bounds $p_L(3) = 6.5536 \times 10^{-4}$ and $p_L(4) = 4.29 \times 10^{-7}$; suppression factors of approximately 6.10×10^2 and 9.31×10^5; successive majority-vote failure probabilities beginning at 0.1736704; and ball-code sizes $n_1 = 5$, $n_2 = 17$, $n_3 = 53$ [8]. Our independent derivations in Section 4 reproduce each of these values from first principles, which serves as a cross-check of both the model and the arithmetic.

A companion QNFO analytic-derivation paper obtains the same ball sizes from the closed form $n_t = 1 + (p+1)\frac{p^t - 1}{p-1}$ and reports passive logical failure probabilities $P_{\text{pass}} = 9.97 \times 10^{-6}$ at $t = 1$ and $P^{(2)}_{\text{pass}} = 6.71 \times 10^{-7}$ at $t = 2$, together with an energy barrier $E_{\text{barrier}}(3) = 8$, a thermal activation bound $p_{\text{thermal}} = 3.35 \times 10^{-4}$, and an estimate that surface-code overhead is $10^3 : 1$ for comparable logical error rates [9]. We cite these values as reported targets in Sections 4.5 and 7 and compare them with our own computations.

A later QNFO discussion emphasizes the epistemic status of the programme: the barrier scaling $E_{\text{barrier}}(d) = 2^d$ is a posited model rather than a theorem derived from ultrametric geometry alone, the majority-vote cascade is a generic classical property rather than a quantum-specific effect, and the quantitative conclusions depend on simulation reports that have not been independently verified [11]. We carry these caveats forward explicitly in Section 6, since they bound the strength of every claim in this paper.

3. Methods

3.1 Tree Geometry and Notation

We consider a regular odd-p Bruhat–Tits tree $\mathcal{T}_p$ rooted at a logical qudit. Each vertex has degree $p+1$: the root has $p+1$ children, each child has $p+1$ further descendants, and so on. The depth d denotes the number of edges from the root to the farthest leaf of the subtree under consideration. A *ball* of radius t around a vertex comprises all vertices at graph distance at most t; the number of physical qudits in such a ball is

$$n_t = 1 + (p+1)\frac{p^t - 1}{p - 1}, \tag{1}$$

derived by summing the geometric series of nodes at each level (Section 4.4).

3.2 Error Model Each physical qudit suffers an independent depolarizing error with probability ε. Errors are assumed to be Pauli-type and uncorrelated across edges, matching the stochastic noise model used in the barrier-analysis simulations of [8]. This is the binary symmetric channel picture of [4] applied per edge. Coherent error mechanisms are outside the model; we return to this restriction in Section 6.

3.3 Passive Barrier The strong triangle inequality of the ultrametric [7] implies that any error chain connecting two vertices at hierarchical distance d must intersect at least one cut of size 2^d. We model this as an energy barrier

$$B(d) = 2^d, \tag{2}$$

so that the probability that an error chain traverses the cut is $\varepsilon^{B(d)}$. This yields the logical error bound

$$p_L(d) = \varepsilon^{B(d)} = \varepsilon^{2^d}. \tag{3}$$

The derivation of Eq. (2) from the inequality is given step by step in Section 4.1; we flag here that [11] classifies this scaling as a posited model, and we treat it accordingly.

3.4 Recursive Majority-Vote Decoder At each non-leaf vertex a majority vote is taken over the error indications of the subtree below it; the vertex outputs the majority value, with ties broken arbitrarily. Two vote widths appear in the source literature: an 8-leaf binary majority and a vote over the $p+1 = 4$ children of the ternary tree. We take the 8-

leaf majority under the convention of [8], whose reported first-level failure probability 0.1736704 ***our term-by-term arithmetic in Section 4.3 reproduces exactly; for completeness we also record that the narrower 4-child vote gives*** $p^{(1)} = 0.1792$ ***at*** $\varepsilon = 0.4$***, a difference of about*** 3% ***that does not affect any conclusion. The failure probability of one 8-leaf vote with leaf error probability*** q ***is***

$$p^{(1)}(q) = \sum_{k=5}^{8} \binom{8}{k} q^k (1-q)^{8-k}, \tag{4}$$

and the recursion $q_{i+1} = p^{(1)}(q_i)$ produces the failure probability after i concatenated levels.

3.5 Parameter Choices and Assumptions We fix $p = 3$ ***(ternary tree) and*** $\varepsilon = 0.4$ ***to match the numerical values reported in [8], so that our independent arithmetic can be checked line by line against a published target. For the prototype proposal in Section 7 we instead assume a calibrated depolarizing channel with*** $\varepsilon = 10^{-3}$***, a value we state as an assumption for the experiment design rather than a measurement. All calculations are analytical; no numerical simulation is performed in this paper.***

4. Analysis

4.1 Barrier Height Derivation

The ultrametric strong triangle inequality [7] states that for any three vertices x, y, z,

$$d(x, y) \leq \max\{d(x, z), d(z, y)\}.$$

Consider a root-to-leaf path of length d. To connect the root to a leaf, an error chain must traverse at least one edge at each hierarchical level, because any shortcut that skips a level would place two vertices at the same ultrametric distance on the path and violate the inequality. Hence the minimal number of edges that must be flipped simultaneously to produce a logical error at the root is bounded by a quantity that doubles with each level:

- Level 0 (root) to level 1 requires crossing at least 2 edges (binary branching of the minimal cut).
- Inductively, each additional level doubles the minimal cut size, giving the recurrence $B(d) = 2\,B(d-1)$ with $B(0) = 1$.
- Solving the recurrence gives

$$B(d) = 2^d.$$

Thus $B(3) = 2^3 = 8$ and $B(4) = 2^4 = 16$, matching the barrier heights reported in [8] and the energy barrier $E_{\text{barrier}}(3) = 8$ reported in [9]. We emphasize, following [11], that this doubling is a model of the cut-counting argument, not a theorem proved from the strong triangle inequality alone; the derivation above assumes that the minimal cut at level ℓ has exactly 2^ℓ edges, which holds for the idealized regular tree but need not hold for hardware with defects or long-range couplings.

4.2 Logical Error Probability and Suppression Factor Substituting Eq. (2) into Eq. (3) with $\varepsilon = 0.4$***:***

$$p_L(3) = 0.4^{2^3} = 0.4^8.$$

Step by step:

$$0.4^2 = 0.16, \qquad 0.4^4 = 0.16^2 = 0.0256, \qquad 0.4^8 = 0.0256^2 = 0.00065536,$$

so

$$p_L(3) = 6.5536 \times 10^{-4},$$

which equals the value listed in [8]. For depth $d = 4$:

$$p_L(4) = 0.4^{2^4} = 0.4^{16} = \left(0.4^8\right)^2 = \left(6.5536 \times 10^{-4}\right)^2 = 4.294967296 \times 10^{-7} \approx 4.29 \times 10^{-7},$$

again matching [8].

The suppression factor relative to the physical error rate is

$$S(d) = \frac{\varepsilon}{p_L(d)}.$$

For $d = 3$:

$$S(3) = \frac{0.4}{6.5536 \times 10^{-4}} = 610.35\ldots \approx 6.10 \times 10^2.$$

For $d = 4$:

$$S(4) = \frac{0.4}{4.294967296 \times 10^{-7}} = 9.3132\ldots \times 10^5 \approx 9.31 \times 10^5.$$

Both values match the suppression factors reported in [8].

4.3 Recursive Majority-Vote Failure Probabilities Using Eq. (4) with leaf error probability $q = \varepsilon = 0.4$:

$$\begin{aligned} p^{(1)} &= \binom{8}{5} 0.4^5 0.6^3 + \binom{8}{6} 0.4^6 0.6^2 + \binom{8}{7} 0.4^7 0.6^1 + \binom{8}{8} 0.4^8 \\ &= 56 \cdot 0.01024 \cdot 0.216 + 28 \cdot 0.004096 \cdot 0.36 \\ &\quad + 8 \cdot 0.0016384 \cdot 0.6 + 1 \cdot 0.00065536 \\ &= 0.12386304 + 0.04128768 + 0.00786432 + 0.00065536 \\ &= 0.1736704. \end{aligned}$$

This reproduces the first-level failure probability reported in [8] exactly. For comparison, the narrower 4-child vote of Section 3.4 gives

$$p_4^{(1)} = \binom{4}{3} 0.4^3 0.6^1 + \binom{4}{4} 0.4^4 = 4 \cdot 0.064 \cdot 0.6 + 0.0256 = 0.1536 + 0.0256 = 0.1792,$$

about 3% higher; we retain the 8-leaf convention throughout.

For the second level we set $q = p^{(1)} = 0.1736704$ in Eq. (4). With $1 - q = 0.8263296$:

$$q^2 = 0.0301614, \quad q^3 = 0.0052399, \quad q^4 = 0.00091010, \quad q^5 = 1.5807 \times 10^{-4},$$

$$q^6 = 2.7452 \times 10^{-5}, \quad q^7 = 4.7677 \times 10^{-6}, \quad q^8 = 8.2784 \times 10^{-7},$$

$$(1-q)^2 = 0.68282, \qquad (1-q)^3 = 0.56424.$$

The four terms are:

$$\begin{aligned} k = 5: &\quad 56 \cdot 1.5807 \times 10^{-4} \cdot 0.56424 = 4.9950 \times 10^{-3}, \\ k = 6: &\quad 28 \cdot 2.7452 \times 10^{-5} \cdot 0.68282 = 5.2485 \times 10^{-4}, \\ k = 7: &\quad 8 \cdot 4.7677 \times 10^{-6} \cdot 0.8263296 = 3.1517 \times 10^{-5}, \\ k = 8: &\quad 8.2784 \times 10^{-7}. \end{aligned}$$

Summing:

$$p^{(2)} = 4.9950 \times 10^{-3} + 5.2485 \times 10^{-4} + 3.1517 \times 10^{-5} + 8.2784 \times 10^{-7} \approx 5.552 \times 10^{-3},$$

consistent with the value 5.5489×10^{-3} reported in [8]; the small difference reflects rounding of the intermediate powers above.

For the third level we set $q = p^{(2)} = 5.5489 \times 10^{-3}$ (using the reported value of [8] as input). With $1 - q = 0.9944511$:

$$q^2 = 3.0790 \times 10^{-5}, \quad q^3 = 1.7086 \times 10^{-7}, \quad q^4 = 9.4809 \times 10^{-10}, \quad q^5 = 5.2610 \times 10^{-12},$$

$$q^6 = 2.9194 \times 10^{-14}, \quad q^7 = 1.6199 \times 10^{-16}, \quad q^8 = 8.988 \times 10^{-19},$$

$$(1-q)^3 = 0.98347, \qquad (1-q)^2 = 0.98894.$$

The dominant term is

$$k = 5: \quad 56 \cdot 5.2610 \times 10^{-12} \cdot 0.98347 = 2.897 \times 10^{-10},$$

with corrections

$$k = 6: \ 28 \cdot 2.9194 \times 10^{-14} \cdot 0.98894 = 8.085 \times 10^{-13}, \qquad k = 7: \ 8 \cdot 1.6199 \times 10^{-16} \cdot 0.9944511 = 1.288 \times 10^{-15}.$$

Summing:

$$p^{(3)} \approx 2.897 \times 10^{-10} + 8.085 \times 10^{-13} + 1.288 \times 10^{-15} \approx 2.905 \times 10^{-10},$$

matching the third-level value listed in [8]. Each concatenation level therefore suppresses the failure probability by roughly two orders of magnitude at this error rate.

4.4 Ball-Code Size Formulas A ball of radius t in a regular odd-p tree contains the root plus all vertices up to depth t. Level 0 contributes 1 node; each subsequent level $\ell = 1, \dots, t$ contributes $(p+1)p^{\ell-1}$ nodes. Summing the geometric series,

$$n_t = 1 + (p+1)\sum_{\ell=1}^{t} p^{\ell-1} = 1 + (p+1)\frac{p^t - 1}{p-1}.$$

Evaluating for $p = 3$:

$$\begin{aligned} n_1 &= 1 + 4 \cdot \frac{3^1 - 1}{3 - 1} = 1 + 4 \cdot \frac{2}{2} = 1 + 4 = 5, \\ n_2 &= 1 + 4 \cdot \frac{3^2 - 1}{3 - 1} = 1 + 4 \cdot \frac{8}{2} = 1 + 16 = 17, \\ n_3 &= 1 + 4 \cdot \frac{3^3 - 1}{3 - 1} = 1 + 4 \cdot \frac{26}{2} = 1 + 52 = 53. \end{aligned}$$

Thus a depth-1 ball uses 5 qudits, depth-2 uses 17, and depth-3 uses 53. These values coincide with the ball-code sizes listed in [8] and derived in [9].

4.5 Energy Barrier and Thermal Activation Using the barrier height $B(d) = 2^d$ we define an energy barrier in arbitrary units,

$$E_{\text{barrier}}(d) = 2^d,$$

which matches the value $E_{\text{barrier}}(3) = 8$ reported in [9]. Assuming a thermal activation model with Boltzmann factor $e^{-E_{\text{barrier}}/k_B T}$ and setting $k_B T = 1$ (unit temperature, an assumption of the model), the thermal error probability is

$$p_{\text{thermal}}(d) = e^{-2^d}.$$

For $d = 3$:

$$p_{\text{thermal}}(3) = e^{-8} = 3.3546 \times 10^{-4} \approx 3.35 \times 10^{-4},$$

$$p_{\text{thermal}}(3) \approx 3.35 \times 10^{-4},$$

which matches the thermal activation bound $p_{\text{thermal}} = 3.35 \times 10^{-4}$ reported in [9]. For completeness, the depth-5 value is

$$p_{\text{thermal}}(5) = e^{-32} \approx 1.27 \times 10^{-14},$$

consistent with the projection $p_L(5) \leq 1.2664 \times 10^{-14}$ labeled in [11]. We stress that this bound holds only under the posited barrier model with no tunnelling or coherent path, as [11] states explicitly; any such path invalidates it entirely.

5. Results

All numbers below are computed in Section 4 from the stated inputs ($p = 3$, $\varepsilon = 0.4$, unit temperature $k_B T = 1$); no simulation or measurement was performed in this paper.

Quantity	Value	Derived in
Barrier height $B(3)$	8	Section 4.1
Barrier height $B(4)$	16	Section 4.1
Logical error bound $p_L(3)$	6.5536×10^{-4}	Section 4.2
Logical error bound $p_L(4)$	$4.294967296 \times 10^{-7} \approx 4.29 \times 10^{-7}$	Section 4.2
Suppression factor $S(3)$	$\approx 6.10 \times 10^{2}$	Section 4.2
Suppression factor $S(4)$	$\approx 9.31 \times 10^{5}$	Section 4.2
First-level 8-leaf vote failure $p^{(1)}$	0.1736704	Section 4.3
Second-level failure $p^{(2)}$	$\approx 5.55 \times 10^{-3}$	Section 4.3
Third-level failure $p^{(3)}$	$\approx 2.905 \times 10^{-10}$	Section 4.3
4-child vote failure $p_4^{(1)}$	0.1792	Section 4.3
Ball-code size n_1 ($p = 3$)	5 qudits	Section 4.4
Ball-code size n_2 ($p = 3$)	17 qudits	Section 4.4
Ball-code size n_3 ($p = 3$)	53 qudits	Section 4.4
Thermal activation bound $p_{\text{thermal}}(3)$	$e^{-8} \approx 3.35 \times 10^{-4}$	Section 4.5
Thermal activation bound $p_{\text{thermal}}(5)$	$e^{-32} \approx 1.27 \times 10^{-14}$	Section 4.5

Each row reproduces, by independent arithmetic, a value reported in the barrier analysis of [8], the analytic-derivation paper of [9], or the corrected analysis of [11]. The agreement covers the barrier heights, the logical error bounds, the suppression factors, the majority-vote cascade, and the ball-code sizes; the only quantity we compute that [9] reports at a different value is the passive logical failure probability, where [9] gives $P_{\text{pass}} = 9.97 \times 10^{-6}$ at $t = 1$ under its own noise assumptions, which differ from the $\varepsilon = 0.4$ convention we adopt for cross-checking, so the two figures are not directly comparable.

6. Discussion

The central result of this paper is conditional: under the posited barrier model $B(d) = 2^d$ and independent depolarizing errors with $\varepsilon = 0.4$, ultrametric tree connectivity suppresses the logical error probability from 6.55×10^{-4} at depth 3 to 4.29×10^{-7} at depth 4, a factor of 9.31×10^{5} relative to the physical error rate. The strength of this result is bounded by four limitations, each of which we state explicitly.

The barrier scaling is a model, not a theorem. As [11] establishes, the cut-counting argument behind $B(d) = 2^d$ counts sites on a level of any rooted binary tree and invokes the strong triangle inequality nowhere; [11] concludes that the barrier $E_{\text{barrier}}(d) = 2^d$ is a posited model rather than a derivation from ultrametric geometry. Our Section 4.1 derivation inherits

this weakness: the inductive step assumes the minimal cut at level ℓ has exactly 2^ℓ edges, which holds for the idealized regular tree but not for hardware with defects, asymmetric couplings, or long-range links. A hardware-realistic tree with a missing edge at level ℓ reduces the local cut size and can dominate the logical error rate.

The majority-vote cascade is classical, not quantum-specific. The recursion $0.1736704 \to 5.55 \times 10^{-3} \to 2.905 \times 10^{-10}$ computed in Section 4.3 is a property of recursive majority voting on any branching structure, as [11] argues: it confirms a decoder-level classical fact, not a mechanism of the tree geometry. The transfer of this classical cascade to quantum errors requires that coherent errors cannot survive majority voting; [8] lists exactly this as a falsification condition, noting that a unitary error acting on a branch that can be reconstructed at the root despite majority-vote decoding would invalidate the classical-to-quantum barrier transfer.

Independence of errors. The depolarizing model of Section 3.2 assumes errors are uncorrelated across edges. Correlated noise—for example a gate fault coupling two neighbouring vertices, the correlated-error limitation stated in [11]—would violate the binomial sums of Section 4.3. If a physical mechanism flips $\lceil p/2 \rceil$ edges simultaneously with non-negligible probability, the barrier exponent is overestimated, again as [8] states in its failure-mode analysis.

Unverified simulation reports. The numerical targets we cross-check against come from the QNFO analyses themselves; [11] states that the single empirical input of the programme—zero logical errors at $p = 0.4$ for $d \geq 3$—is an unverified report, that no cited work documents the simulation, its trial count, or its decoder, and that even for $N = 10^5$ trials the reported zero-error result bounds the logical error probability only at 2.9957×10^{-5}. Our arithmetic reproduces the *reported* values, which validates internal consistency of the model chain but not the model's physical validity.

What would falsify the main result. The main result—the exponential barrier $p_L(d) = \varepsilon^{2^d}$—would be falsified by any of the three observations [8] enumerates: (i) measuring $p_L(d)$ for increasing d and finding a depth-independent plateau, contradicting Eq. (3); (ii) demonstrating a coherent error path that survives majority voting; or (iii) showing a correlated mechanism that flips $\lceil p/2 \rceil$ edges simultaneously with non-negligible probability. A fourth falsifier, specific to our arithmetic, is a discrepancy between the closed-form cascade and a hardware measurement at matched noise: if a five-atom prototype at $\varepsilon = 10^{-3}$ showed a logical error probability above the threshold derived in Section 7, the confinement model would be refuted at the smallest nontrivial depth.

Limitations list.

1. *Barrier model*: $B(d) = 2^d$ is posited, not derived from the strong triangle inequality alone [11]; sensitivity is total, since every downstream number is a power of ε with exponent 2^d.
2. *Independence*: all binomial sums assume uncorrelated per-edge errors; a single correlated mechanism at rate ε_c adds a term of order ε_c that can dominate ε^{2^d} at depth ≥ 3.
3. *Coherent errors*: outside the model entirely; the classical-to-quantum transfer is unproven [8].
4. *Ties in majority vote*: ties are broken arbitrarily in Section 3.4; for the 8-leaf vote ties cannot occur, but for the 4-child vote a 2–2 tie occurs with probability $\binom{4}{2} q^2 (1-q)^2$, which at $q = 0.4$ is $6 \cdot 0.16 \cdot 0.36 = 0.3456$, so the 4-child variant requires a tie-breaking rule whose effect we have not modeled.
5. *Cross-check targets*: the values in [8] and [9] are unverified reports [11]; agreement with them checks arithmetic, not physics.
6. *Thermal model*: $k_B T = 1$ is an assumption with no calibrated physical unit; the bound $p_{\text{thermal}}(d) = e^{-2^d}$ holds only with no tunnelling [11].

Open questions. Whether a quantum encoder exists whose logical states are eigenstates of the subtree confinement operator remains open; [11] notes that a related distance result $d = 2t + 1$ is conditional on an unconstructed branch-permutation-symmetric encoder. Whether the ultrametric barrier survives under finite-temperature coherent dynamics, rather than the stochastic channel of Section 3.2, is likewise open.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The prospective users are quantum hardware architects on platforms with reconfigurable connectivity—neutral-atom tweezer arrays and ion-trap shuttling architectures—whose interconnect topology is a design parameter rather than a fixed constraint. If the barrier model holds, the design change is to route the connectivity graph as a regular odd-p tree instead of a 2D nearest-neighbour lattice, so that a logical qudit is protected by a ball code of 5 to 53 qudits (Section 4.4) rather than the 10^2 to 10^4 physical qubits of surface-code layouts. [9] frames exactly this trade, estimating a surface-code overhead of $10^3 : 1$ against $5 : 1$ for the ultrametric ball code, and reports that under its assumptions the passive scheme matches surface-code logical errors using five-fold fewer physical qudits. We treat the $10^3 : 1$ figure as an internal estimate of [9], not as a derived result of this paper. Code designers and decoder engineers are secondary users: the recursive majority-vote decoder of Section 3.4 requires no syndrome-extraction schedule, replacing real-time classical decoding with a fixed feed-forward vote tree.

(b) Readiness level. This work is purely analytical and has not been peer reviewed; under the NASA Technology Readiness Level definitions it sits at TRL 1–2: basic principles observed and reported, with the barrier model formulated and the arithmetic cross-checked, but no experimental proof of concept. Analytical work cannot exceed TRL 3, and no manufacturing readiness is claimed. The claims of [9] that rest on simulation reports remain, per [11], unverified.

(c) Smallest next experiment. The minimal falsification test, following the prototype design of [9], is a five-^{87}Rb-atom optical-tweezer array arranged as a $p = 3$ tree of depth $t = 1$. The proposed protocol would: (1) initialize each atom in a chosen hyperfine level encoding a qudit; (2) apply a controlled depolarizing channel with assumed $\varepsilon = 10^{-3}$, calibrated by laser-noise injection (this ε is an assumption of the design, not a measurement); (3) perform a majority-vote readout of the root ball after one protection epoch. [9] states the decision rule: if the observed logical error probability exceeds 2×10^{-5}, the ultrametric confinement model would be falsified. A complementary in-silico test is a Monte Carlo comparison of concatenated decoding on tree versus 2D connectivity graphs under matched local stochastic noise, where a $\geq 2\times$ threshold improvement confirms and a null result falsifies the geometric claim; this simulation has not been performed in this paper and remains a proposed de-risking step.

(d) Binding constraints. The constraints that would bind first, in order: (i) *connectivity fidelity*: the barrier requires the regular tree to be defect-free, since a single missing edge at level ℓ reduces the minimal cut below 2^ℓ (Section 4.1); (ii) *correlated noise*: any two-qudit gate fault couples neighbouring vertices and breaks the independence assumption of Section 4.3, so gate crosstalk must be characterized before the model applies; (iii) *readout fidelity*: the majority vote consumes one qudit readout per leaf, so readout error adds directly to the leaf error probability q and must satisfy $q_{\text{readout}} \ll q_{\text{gate}}$; (iv) *coherence time*: [9] reports that QND monitoring imposes a coherence-time ceiling $\tau_{\text{H}} \sim 2^L/(JL)$ that limits practical depth to $t \leq 1$ on current hardware, which would cap the achievable barrier at $B(1) = 2$ until coherence improves. Materials, tolerance, and regulatory constraints are open questions that this paper does not quantify.

(e) Order-of-magnitude bounds. As a labeled projection with stated assumptions: if the depth-3 ball code (53 qudits, Section 4.4) achieved the suppression $S(3) \approx 6.10 \times 10^2$ at $\varepsilon = 0.4$ (Section 4.2), the qudit cost per logical qudit would be 53 against the 10^2–10^4 range of surface-code layouts stated in Section 1, representing a factor of about 1.9 relative to the lower bound 10^2 and about 1.9×10^2 relative to the upper bound 10^4. This projection assumes the barrier model, uncorrelated depolarizing noise, and defect-free tree connectivity; it is not a hardware cost estimate, and no fabrication, power, or dollar figure is derivable from the grounding.

(f) Risks and misuse paths. The principal risk is overclaiming: presenting the arithmetic agreement of Section 5 with [8] and [9] as experimental validation, when [11] shows the underlying simulation reports are unverified and the barrier is a posited model. A secondary risk is that the epistemically critical framing of [10], that results can be retroactively justified because the numbers fit, applies here if the barrier model is adopted without the falsification test of (c); the five-atom experiment is the safeguard against exactly this failure mode.

8. Conclusion

We analyzed a conjectured passive fault-tolerance mechanism in which logical qudits are encoded on a regular odd-p Bruhat–Tits tree and the ultrametric strong triangle inequality confines correlated errors within subtrees. Under the posited barrier model $B(d) = 2^d$ and independent depolarizing noise with $\varepsilon = 0.4$, we derived by explicit arithmetic the logical error bounds $p_L(3) = 6.5536 \times 10^{-4}$ and $p_L(4) = 4.29 \times 10^{-7}$, suppression factors $S(3) \approx 6.10 \times 10^2$ and $S(4) \approx 9.31 \times 10^5$, the recursive majority-vote cascade $0.1736704 \to 5.55 \times 10^{-3} \to 2.905 \times 10^{-10}$, ball-code sizes 5, 17, and 53 qudits at depths 1–3, and the thermal activation bound $p_{\text{thermal}}(3) = e^{-8} \approx 3.35 \times 10^{-4}$. Each value reproduces, by independent computation, a target reported in the prior barrier analyses [8], [9], cross-checked against the corrected analysis of [11].

The strength of these results is bounded by the caveats we carry forward: the barrier scaling is a posited model rather than a theorem of ultrametric geometry [11], the majority-vote cascade is a generic classical property of any branching structure [11], the error-independence assumption excludes correlated and coherent mechanisms [8], and the cross-check targets are unverified simulation reports [11]. The result therefore stands as a well-defined, internally consistent model with a sharp falsification criterion: a five-atom depth-1 prototype at assumed $\varepsilon = 10^{-3}$ whose logical error probability exceeding 2×10^{-5} would refute the confinement model [9]. If the model survives that test and the Monte Carlo tree-versus-2D threshold comparison, ultrametric connectivity would make fault tolerance partly a passive geometric property of hardware, reducing logical-qubit overhead from the 10^2–10^4 range toward tens of qudits; if it fails, the programme reduces to classical recursive majority voting, which needs no ultrametric geometry. The analysis also connects to adjacent hierarchical-geometry problems—holographic codes, spin-glass ultrametricity, and p-adic condensed-matter models—in which the same strong triangle inequality governs how information and disturbance propagate across scales.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.55.0). 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 14 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] Continuous-time quantum error correction. arXiv:1311.2485v2. https://arxiv.org/abs/1311.2485v2 [2] Entanglement-Assisted Quantum Error-Correcting Codes. arXiv:1610.04013v1. https://arxiv.org/abs/1610.04013v1 [3] Quantum Error Correction in Quaternionic Hilbert Spaces. arXiv:2504.19833v1. https://arxiv.org/abs/2504.19833v1 [4] An Introduction to Error-Correcting Codes: From Classical to Quantum. arXiv:quant-ph/0602157v1. https://arxiv.org/abs/quant-ph/0602157v1

[5] Quantum Computing and Error Correction. arXiv:quant-ph/0304016v2. https://arxiv.org/abs/quant-ph/0304016v2 [6] Quantum Error Correction. arXiv:1910.03672v1. https://arxiv.org/abs/1910.03672v1 [7] An embedding, an extension, and an interpolation of ultrametrics. arXiv:2008.10209v2. https://arxiv.org/abs/2008.10209v2 [8] https://papers.qnfo.org/papers/ultrametric-passive-error-correction-on-p-ary-bruhat-tits-trees-barrier-analysis. QNFO: Ultrametric Passive Error Correction on *p*-ary Bruhat–Tits Trees: Barrier Analysis and Recursive Thresholds. [9] https://papers.qnfo.org/papers/ultrametric-passive-fault-tolerance-on-bruhat-tits-trees-analytic-derivation-and. QNFO: Ultrametric Passive Fault Tolerance on Bruhat–Tits Trees: Analytic Derivation and Numerical Projections. [10] https://papers.qnfo.org/papers/quantum-fraud. QNFO: Quantum Fraud. [11] https://papers.qnfo.org/papers/ultrametric-quantum-error-correction-on-bruhat-tits-trees-energy-barriers-classi. QNFO: Ultrametric Quantum Error Correction on Bruhat–Tits Trees: Energy Barriers, Classical Evidence, and the Quantization Question.