

V1.0.0

Passive Fault Tolerance via Bruhat–Tits Tree Codes: Exact Ball-Code Error Rates, Energy Barriers, and the Cooling Wall of Active Quantum Error Correction

Rowan Brad Quni-Gudzin

QNFO-ULA

ABSTRACT

Large-scale quantum computation is gated by active error correction: stabilizer codes such as the surface code require physical-to-logical overheads of 10^2 – 10^4 , continuous syndrome measurement, and millikelvin refrigeration whose thermodynamic cost grows as the operating temperature falls. A proposed alternative encodes logical qudits on balls of the Bruhat–Tits tree T_{p+1} , so that the ultrametric strong triangle inequality confines errors to single branches and correction is passive. The gap addressed here is quantitative: prior statements of the scheme assert a depolarizing threshold near 75% and removal of the cooling wall, but neither figure has been derived. We formalize the ball-code model, derive the exact sizes $n_t = 1 + (p+1)(p^t - 1)/(p - 1)$, and compute with full arithmetic the passive logical failure probabilities $P_{\text{pass}}(1) = 9.980 \times 10^{-6}$ per epoch for the $p = 3$, $t = 1$ block (5 qudits, distance 3) and $P_{\text{pass}}(2) = 6.729 \times 10^{-7}$ for the $t = 2$ block (17 qudits, distance 5) at $\varepsilon = 10^{-3}$, suppression factors of 100.2 and 1486. Under the conjectured barrier $E_{\text{barrier}}(d) = 2^d$, thermal activation gives $p_L \leq 3.355 \times 10^{-4}$ at depth 3. The Carnot coefficient-of-performance ratio between 4 K and 10 mK operation is 405. The 75% threshold remains an open conjecture; encoder existence is the binding constraint.

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-Gudzin, R. B. (n.d.). Passive Fault Tolerance via Bruhat–Tits Tree Codes: Exact Ball-Code Error Rates, Energy Barriers, and the Cooling Wall of Active Quantum Error Correction. QNFO Preprint v1.0.0.

Abstract

Large-scale quantum computation is gated by active error correction: stabilizer codes such as the surface code require physical-to-logical overheads of 10^2 – 10^4 , continuous syndrome measurement, and millikelvin refrigeration whose thermodynamic cost grows as the operating temperature falls. A proposed alternative encodes logical qudits on balls of the Bruhat–Tits tree T_{p+1} , so that the ultrametric strong triangle inequality confines errors to single branches and correction is passive. The gap addressed here is quantitative: prior statements of the scheme assert a depolarizing threshold near 75% and removal of the cooling wall, but neither figure has been derived. We formalize the ball-code model, derive the exact sizes $n_t = 1 + (p+1)(p^t - 1)/(p - 1)$, and compute with full arithmetic the passive logical failure probabilities $P_{\text{pass}}(1) = 9.980 \times 10^{-6}$ per epoch for the $p = 3$, $t = 1$ block (5 qudits, distance 3) and $P_{\text{pass}}(2) = 6.729 \times 10^{-7}$ for the $t = 2$ block (17 qudits, distance 5) at $\varepsilon = 10^{-3}$, suppression factors of 100.2 and 1486. Under the conjectured barrier $E_{\text{barrier}}(d) = 2^d$, thermal activation gives $p_L \leq 3.355 \times 10^{-4}$ at depth 3. The Carnot coefficient-of-performance ratio between 4 K and 10 mK operation is 405. The 75% threshold remains an open conjecture; encoder existence is the binding constraint.

1. Introduction

Quantum processors today face a scaling wall that is thermodynamic as much as it is architectural. Fault tolerance in every leading architecture — surface codes, color codes, and other stabilizer schemes — is active: errors are detected by repeatedly measuring syndrome operators and corrected by classical decoding, with physical-to-logical qubit overheads that published estimates place between 10^2 and 10^4 [10] and that the Alpha Pi project analysis summarizes as roughly 1000:1 for surface

codes [11]. Syndrome extraction must run continuously, at low latency, on hardware held near 10–20 mK, because the measurement and control layer itself must be error-free enough not to overwhelm the code [4]. The consequence is that the refrigeration plant, not the qubit count alone, becomes a dominant cost: cooling power scales inversely with the Carnot coefficient of performance, which collapses as the operating temperature falls, and the Alpha Pi analysis argues that active correction additionally hits a signal-processing ceiling beyond which no classical filtering removes fundamental quantum back-action [11]. Cloud quantum-service providers feel this directly as higher subscription cost and shorter runnable circuits; hardware teams feel it as dilution-refrigerator capacity, wiring, and heat-load budgets.

A design that suppressed logical errors passively — by geometry rather than by measurement — would change three decisions at once. It would remove the syndrome-measurement layer and its ancilla overhead, which traditional fault-tolerant schemes require in numbers proportional to the maximum stabilizer weight [6]. It would relax the refrigeration requirement from the millikelvin regime toward the 4 K regime reachable by pulse-tube cryocoolers, where the Carnot coefficient of performance is roughly 405 times larger than at 10 mK (computed in Section 4.6). And it would convert the scalability question from an engineering race against decoherence into a checkable geometric claim. The proposal under examination, due to the ultrametric quantum error correction program [9, 10, 11, 12], asserts exactly this: encoding on the balls of a Bruhat–Tits tree creates an energy barrier $E_{\text{barrier}}(d) = 2^d$ against logical error, confines error propagation to single branches by the ultrametric strong triangle inequality, and — in its strongest form — admits depolarizing error thresholds near 75%, far above the roughly 1% figure the proposal attributes to surface codes.

The gap this paper closes is quantitative and critical rather than promotional. The 75% threshold and the full removal of the cooling wall are assertions; what can be derived today is narrower and still useful: exact code parameters, exact passive failure probabilities under stated noise assumptions, the barrier scaling and its thermal-activation consequences, and the thermodynamic ratio that quantifies how much of the cooling wall passive operation would remove. We derive these numbers with fully shown arithmetic, we state which claims remain conjectural, and we carry the falsification criteria of the source program [9, 10, 12] into a form a laboratory test could reject.

The contributions are:

- Derivation of the exact ball-code sizes $n_t = 1 + (p+1)(p^t - 1)/(p-1)$ for T_{p+1} , giving $n_1 = 5$, $n_2 = 17$, $n_3 = 53$ for $p = 3$ and $n_1 = 4$, $n_2 = 10$, $n_3 = 22$ for $p = 2$ (Section 4.1).
- Exact computation of the passive logical failure probability $P_{\text{pass}}(1) = 9.980 \times 10^{-6}$ per protection epoch for the $p = 3$, $t = 1$ block at $\varepsilon = 10^{-3}$, a suppression factor of 100.2, and $P_{\text{pass}}(2) = 6.729 \times 10^{-7}$ for the $t = 2$ block, a suppression factor of 1486 (Sections 4.2–4.4, results in Section 5).
- Derivation of the thermal-activation bounds $p_L(3) \leq e^{-8} = 3.355 \times 10^{-4}$ and $p_L(5) \leq e^{-32} = 1.266 \times 10^{-14}$ under the conjectured barrier $E_{\text{barrier}}(d) = 2^d$ with $k_B T = 1$ (Section 4.5).
- Computation of the Carnot coefficient-of-performance ratio 405 between 4 K and 10 mK operation, quantifying the fraction of the cooling wall that passive correction removes under fixed heat load (Section 4.6).
- A critical statement of the quantization gap: the 75% depolarizing threshold and the quantum confinement of superposed error paths (Hypothesis Q of [9]) remain unproved, with explicit falsification criteria (Sections 6 and 7).

2. Background and Related Work

The geometric setting of this paper is the Bruhat–Tits tree, and the arithmetic of its subgraphs is well developed in the number-theoretic literature. Work on the missing branches of the Bruhat–Tits tree studies the set $S(\Omega)$ of maximal orders in the local matrix algebra $\mathfrak{A} = \mathbb{M}_2(k)$ containing a given order Ω , a set realized as a sub-tree of the Bruhat–Tits tree called the branch of the order, and uses these branches to study the selectivity problem and normalizers of Eichler orders [1]. That work fixes the graph conventions this paper adopts — vertices with a symmetric neighborhood relation, full subgraphs, and vertex valency — and its proof machinery walks ghost paths and stems inside the tree, computing distances between branch

structures case by case [1]. The relevance here is foundational: the branch decomposition of [1] is the same tree geometry whose nesting and disjointness properties the passive code exploits, and the valuation-theoretic classification of ramified units and uniformizers in [1] underlies the qudit alphabet structure used in Section 3.

Against this geometric proposal stands the entire active fault-tolerance paradigm. Homological product codes present a method for implementing a universal set of fault-tolerant logical gates by fault-tolerantly mapping between different encoded representations of a given logical state, so that different classes of transversal gates of the underlying codes become available and the usual gate restrictions of any single encoding are circumvented [3]. The construction is explicitly active: the logical state is held in a code whose representation must be switched by an actively verified fault-tolerant procedure, with syndrome information extracted and processed at each mapping step. The passive scheme of this paper removes the measurement layer entirely, so the comparison in Section 6 is between the gate flexibility that homological products buy [3] and the syndrome-free operation that the tree geometry promises.

Recent work has attacked the overhead of active schemes from the protocol side. Low-overhead transversal fault tolerance for universal quantum computation demonstrates that logical operations can be transversal and interleaved with a single syndrome-extraction round instead of d rounds, using partial decoding and mid-circuit commitment, with circuit-level simulations of a $|\overline{S}\rangle$ distillation factory based on the $[[7, 1, 3]]$ Steane code at fixed circuit noise $p = 0.1\%$ [4]. The same work notes that transversal operations and correlated decoding were recently found highly effective in experiments with reconfigurable neutral-atom arrays [4]. This matters for the present proposal in two ways: it shows the active-overhead baseline is a moving target that any passive scheme must beat numerically, and it identifies neutral-atom tweezer arrays as the natural platform for the minimal experiment proposed in Section 7.

On the ancilla side, the flag technique achieves fault-tolerant error correction on capped color codes with few ancillas by detecting high-weight errors arising from a few faults using a small number of flag ancillas, and the work observes that distinguishability of fault sets depends on the number of flag ancillas, gate ordering, and the choice of measured generators [6]. The same introduction records that the Shor error-correction scheme requires as many ancillas as the maximum weight of the stabilizer generators [6]. The passive scheme's selling point is the limit of this trajectory: zero syndrome measurements and therefore zero ancillas, which Section 4 prices into the heat budget.

The broader fault-tolerance literature supplies both analogies and methodological discipline. The orthogonal fault tolerance model was proposed to manage and reason about multiple fault tolerance mechanisms that co-exist in dynamically adaptive systems, where the interaction of co-evolving mechanisms is the key challenge [2]; a passive geometric layer coexisting with any residual active layer is precisely such a multi-mechanism design, and feature interaction between the layers is a risk Section 6 takes seriously. In cloud systems, Byzantine faults are highlighted as especially dangerous because they often go undetected at the initial stage and can propagate to other virtual machines before detection, keeping mission-critical applications off the cloud [5] — the same propagation-without-detection failure mode that the ultrametric confinement claim is designed to exclude. Decentralized validation for non-malicious arbitrary fault tolerance in Paxos hardened an active-replication framework and validated it by fault injection: 100 message-fault injections and 100 application-fault injections were each detected with zero errors, and the summed coverage of protocol-deviation faults reached 97.6% [7]. That injection methodology — inject faults, count escapes — is the template for the falsification experiment of Section 7. Finally, SignSGD demonstrates fault tolerance to blind and Byzantine adversaries in distributed learning: the aggregation of signed gradients maintains accuracy as the number of blind adversaries increases, with a theoretical bound verified in practice on logistic regression and feed-forward networks [8]. This is the closest classical analogue of the passive mechanism: robustness conferred by an aggregation structure (there, sign majority; here, tree-ball majority) rather than by detection and retransmission.

The immediate predecessors of this paper are the four documents of the ultrametric program itself. The energy-barrier paper formalizes the conjecture $E_{\text{barrier}}(d) = 2^d$ in tree depth d , computes the iterated majority-vote suppression of $p = 0.4$ through 0.1736704, 0.0055485, and 2.905×10^{-10} over three levels, bounds a reported zero-error simulation at $p_L \leq 2.9957 \times 10^{-5}$ for $N = 10^5$ trials, and states Hypothesis Q — whether the strong triangle inequality confines superposed and entangled error paths — with explicit falsification criteria, while flagging that the zero-error report is unverified

and that tunnelling through the 2^d -site cut would restore a depth-independent error channel [9]. The analytic-derivation paper assumes a physical-to-logical overhead of 10^2 – 10^4 for active codes, derives compact ball codes ($n_1 = 5$ for $p = 3$, $t = 1$), reports a passive logical failure probability of 5.992003×10^{-6} at $\varepsilon = 10^{-3}$ under its decoder model, and states the 75% threshold conjecture and the $\sim 1\%$ surface-code comparison [10]. The Alpha Pi project document supplies the systems-level framing: the $\sim 1000:1$ surface-code overhead figure, the argument that the refrigeration plant dominates cost through Carnot-limited cooling power, and the signal-processing-ceiling argument against classical filtering of quantum back-action [11]. Finally, the ultrametric-geometry paper proves the nesting/disjointness lemma for equal-radius balls in the p -adic norm and derives the exact combinatorial identity for the ball size n_t used in Section 3 [12]; these two results are the mathematical backbone of the present analysis.

3. Methods

3.1. Tree-code model

We consider the homogeneous Bruhat–Tits tree T_{p+1} associated with the p -adic field $\mathbb{Q}_p$, with the graph conventions of [1]: undirected edges, symmetric neighborhood relation, vertex valency $p + 1$. Each vertex hosts a p -dimensional qudit. A ball of radius t centred at a root vertex consists of all vertices at graph distance $\leq t$; the ball is a finite induced subgraph used as a logical block.

Assumption A1 (noise model). Depolarising noise acts independently on each physical qudit per protection epoch, with error probability ε . Correlated errors across vertices are excluded; Section 6 treats this as a limitation.

Assumption A2 (encoder and decoder). The encoder maps a logical state to a symmetric superposition over all vertices of the ball, and a majority-vote decoder returns the root's value. Under this model the code corrects any error pattern affecting at most t vertices and fails when $t + 1$ or more vertices are corrupted. This is the same majority-vote decoder family for which [9] and [10] report suppression figures; we take it as the working convention throughout, and we note in Section 4.3 where it disagrees with the decoder model of [10].

3.2. Code parameters

The number of vertices in a radius- t ball for branching factor $p + 1$ is

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

Derivation. The root contributes 1 vertex. Layer k ($1 \leq k \leq t$) adds $(p+1)p^{k-1}$ new vertices, so $n_t = 1 + (p+1) \sum_{k=0}^{t-1} p^k = 1 + (p+1) \frac{p^t - 1}{p - 1}$, which is (1). This identity is derived in [12].

The code distance — the minimum number of vertex errors that can cause a logical error — is

$$d = 2t + 1, \quad (2)$$

because the majority-vote decoder corrects up to t erroneous vertices and fails at $t + 1$.

3.3. Logical failure probability

For a ball of size n_t , the decoder fails if at least $t + 1$ vertices are erroneous. Under Assumption A1 the logical failure probability is the binomial tail

$$P_{\text{pass}}(t) = \sum_{k=t+1}^{n_t} \binom{n_t}{k} \varepsilon^k (1 - \varepsilon)^{n_t - k}. \quad (3)$$

We adopt the convention that the tail is summed exactly over all k ; where a leading-term truncation is used we state the truncation explicitly and give both values.

3.4. Suppression factor and thermodynamic quantities

The suppression factor is the ratio of physical to logical error probability,

$$S(t) = \frac{\varepsilon}{P_{\text{pass}}(t)}. \quad (4)$$

For the thermodynamic comparison we use the Carnot coefficient of performance for an ideal refrigerator removing heat from a cold reservoir at T_c and rejecting to T_h :

$$\text{COP}_{\text{Carnot}} = \frac{T_c}{T_h - T_c}. \quad (5)$$

Under the conjectured barrier of [9], the thermal-activation bound on the logical error rate is

$$p_L(d) \leq e^{-E_{\text{barrier}}(d)/k_B T} = e^{-2^d} \quad \text{with } k_B T = 1. \quad (6)$$

4. Analysis

All numerical inputs are either derived here, cited to [9, 10, 11, 12], or stated as explicit assumptions.

4.1. Ball sizes

For $p = 3$, $t = 1$, equation (1) gives

$$n_1 = 1 + \frac{4(3^1 - 1)}{3 - 1} = 1 + \frac{4 \cdot 2}{2} = 1 + 4 = 5,$$

matching the value derived in [10] and [12]. For $p = 3$, $t = 2$:

$$n_2 = 1 + \frac{4(9 - 1)}{2} = 1 + \frac{32}{2} = 1 + 16 = 17.$$

For $p = 3$, $t = 3$:

$$n_3 = 1 + \frac{4(27 - 1)}{2} = 1 + \frac{104}{2} = 1 + 52 = 53.$$

For $p = 2$: $n_1 = 1 + \frac{3(2-1)}{1} = 4$; $n_2 = 1 + \frac{3(4-1)}{1} = 10$; $n_3 = 1 + \frac{3(8-1)}{1} = 22$.

4.2. Logical failure for $t = 1$, $p = 3$

Distance $d = 3$ from (2); the decoder fails if ≥ 2 of the $n_1 = 5$ vertices are erroneous. With $\varepsilon = 10^{-3}$ (assumption; the same value is used in [9, 10]), equation (3) gives four terms:

$$P_{\text{pass}}(1) = \binom{5}{2} \varepsilon^2 (1 - \varepsilon)^3 + \binom{5}{3} \varepsilon^3 (1 - \varepsilon)^2 + \binom{5}{4} \varepsilon^4 (1 - \varepsilon) + \varepsilon^5.$$

Term 1: $\binom{5}{2} = 10$, $\varepsilon^2 = 10^{-6}$, $(1 - \varepsilon)^3 = 0.999^3 = 0.997002999$, so

$$10 \times 10^{-6} \times 0.997002999 = 9.97002999 \times 10^{-6}.$$

Term 2: $\binom{5}{3} = 10$, $\varepsilon^3 = 10^{-9}$, $(1 - \varepsilon)^2 = 0.998001$, so

$$10 \times 10^{-9} \times 0.998001 = 9.98001 \times 10^{-9}.$$

Term 3: $\binom{5}{4} = 5$, $\varepsilon^4 = 10^{-12}$, $(1 - \varepsilon) = 0.999$, so

$$5 \times 10^{-12} \times 0.999 = 4.995 \times 10^{-12}.$$

Term 4: $\varepsilon^5 = 10^{-15}$.

Summing:

$$P_{\text{pass}}(1) = 9.97002999 \times 10^{-6} + 9.98001 \times 10^{-9} + 4.995 \times 10^{-12} + 10^{-15} = 9.980014995 \times 10^{-6} \approx 9.980 \times 10^{-6}.$$

The leading term alone gives 9.970×10^{-6} ; the weight-3 term adds 0.100%, so the exact tail and the leading-term truncation agree to three significant figures at this ε .

4.3. A note on the decoder-model divergence

The analytic-derivation paper reports $P_{\text{pass}} = 5.992003 \times 10^{-6}$ and a suppression factor of 166.9 for the same $p = 3$, $t = 1$, $\varepsilon = 10^{-3}$ block [10], while the ultrametric-geometry paper reports $P_{\text{pass}} \approx 9.98 \times 10^{-6}$ and a suppression factor of ≈ 100.2 [12]. The two predecessor documents use different decoder models for the same ball code. We take the binomial-tail convention of equation (3), under which the decoder fails whenever $t + 1$ or more of the n_t vertices are erroneous, because it is the convention whose inputs are fully stated in [12] (independent per-vertex depolarising noise, majority-vote readout of the root) and it is the conservative choice: it yields the larger failure probability of the two. Under this convention the [10] figure of 5.992003×10^{-6} is not reproduced; we treat the [10] number as belonging to a different, not fully stated decoder model and carry the conservative value 9.980×10^{-6} through the rest of the paper. The same divergence appears at $t = 2$: [10] reports $P_{\text{pass}}^{(2)} = 3.997 \times 10^{-9}$, while [12] reports $\approx 6.73 \times 10^{-7}$; we again take the larger, binomial-tail value.

4.4. Logical failure for $t = 2$, $p = 3$, and suppression factors

Distance $d = 5$ from (2); the decoder fails if ≥ 3 of the $n_2 = 17$ vertices are erroneous. The exact tail has fifteen terms; we compute the first three and bound the remainder.

Term 1 (weight 3): $\binom{17}{3} = 680$, $\varepsilon^3 = 10^{-9}$, $(1 - \varepsilon)^{14} = 0.999^{14}$. Now $\ln(0.999) = -1.00050033 \times 10^{-3}$, so $14 \ln(0.999) = -1.40070047 \times 10^{-2}$ and $0.999^{14} = e^{-0.0140070047} = 0.98608845$. Thus

$$680 \times 10^{-9} \times 0.98608845 = 6.705401 \times 10^{-7}.$$

Term 2 (weight 4): $\binom{17}{4} = 2380$, $\varepsilon^4 = 10^{-12}$, $(1 - \varepsilon)^{13} = e^{-0.0130065043} = 0.98707553$. Thus

$$2380 \times 10^{-12} \times 0.98707553 = 2.349240 \times 10^{-9}.$$

Term 3 (weight 5): $\binom{17}{5} = 6188$, $\varepsilon^5 = 10^{-15}$, $(1 - \varepsilon)^{12} = e^{-0.0120060040} = 0.98806262$. Thus

$$6188 \times 10^{-15} \times 0.98806262 = 6.114020 \times 10^{-12}.$$

The remaining ten terms are bounded by the geometric estimate $\binom{17}{6} \varepsilon^6 (1 - \varepsilon)^{11} \times (1 + 10^{-3} + \dots) \approx 12376 \times 10^{-18} \times 0.989 \times 1.001 \approx 1.225 \times 10^{-14}$, negligible at the reported precision. Summing:

$$P_{\text{pass}}(2) = 6.705401 \times 10^{-7} + 2.349240 \times 10^{-9} + 6.114020 \times 10^{-12} + \mathcal{O}(10^{-14}) = 6.728954 \times 10^{-7} \approx 6.729 \times 10^{-7}.$$

The weight-4 term contributes 0.35% of the total, matching the correction size reported in [12].

Suppression factors from equation (4):

$$S(1) = \frac{10^{-3}}{9.980014995 \times 10^{-6}} = 100.2005 \approx 100.2,$$

$$S(2) = \frac{10^{-3}}{6.728954 \times 10^{-7}} = 1486.17 \approx 1486.$$

The step from $t = 1$ to $t = 2$ buys a further factor $S(2)/S(1) = 1486.17/100.2005 = 14.83$ at the cost of 12 additional qudits ($n_2 = 17$ versus $n_1 = 5$).

4.5. Thermal-activation bounds under the conjectured barrier

The energy-barrier paper conjectures $E_{\text{barrier}}(d) = 2^d$ in tree depth d and derives the thermal-activation form $p_L(d) \leq e^{-E_{\text{barrier}}(d)/k_B T}$ under the assumption $k_B T = 1$ and no tunnelling [9]. These are labeled projections in [9], and we carry them as such here. For $d = 3$: $E_{\text{barrier}}(3) = 2^3 = 8$, so

$$p_L(3) \leq e^{-8} = \frac{1}{e^8} = \frac{1}{2980.958} = 3.35463 \times 10^{-4} \approx 3.355 \times 10^{-4}.$$

For $d = 5$: $E_{\text{barrier}}(5) = 2^5 = 32$, so

$$p_L(5) \leq e^{-32} = \frac{1}{e^{32}} = \frac{1}{7.896296 \times 10^{13}} = 1.26642 \times 10^{-14} \approx 1.266 \times 10^{-14}.$$

Both bounds are conditional on the barrier conjecture, which [9] flags as unproved for quantum dynamics and which would be invalidated entirely by resonant tunnelling through the 2^d -site cut. Note that $p_L(3) = 3.355 \times 10^{-4}$ is larger than the binomial-tail passive failure $P_{\text{pass}}(1) = 9.980 \times 10^{-6}$ at $\varepsilon = 10^{-3}$: at this noise level the combinatorial suppression of the majority-vote decoder, not the thermal barrier, is the operative protection mechanism, and the barrier bound becomes relevant only at temperatures well below the barrier scale.

4.6. Carnot ratio between 4 K and 10 mK operation

We assume heat is rejected to a hot reservoir at $T_h = 300$ K (room-temperature stage), the standard dilution-refrigerator configuration. From equation (5):

At $T_c = 4$ K:

$$\text{COP}_{4\text{K}} = \frac{4}{300 - 4} = \frac{4}{296} = 1.35135 \times 10^{-2}.$$

At $T_c = 10$ mK = 0.01 K:

$$\text{COP}_{10\text{mK}} = \frac{0.01}{300 - 0.01} = \frac{0.01}{299.99} = 3.33444 \times 10^{-5}.$$

The ratio of cooling power available at fixed input work W (since $Q_c = \text{COP} \times W$) is

$$\frac{\text{COP}_{4\text{K}}}{\text{COP}_{10\text{mK}}} = \frac{1.35135 \times 10^{-2}}{3.33444 \times 10^{-5}} = 405.3 \approx 405.$$

Equivalently, removing the same heat load at 4 K requires 405 times less input work than at 10 mK, in the ideal-Carnot limit. This is the quantitative content of the claim that passive operation relaxes the refrigeration requirement from the millikelvin regime toward the 4 K regime: the cooling wall is a factor-405 thermodynamic penalty, not an absolute prohibition. Real refrigerators fall far short of Carnot, so 405 is an upper bound on the achievable ratio; the direction of the effect is robust because both real efficiencies degrade as T_c falls.

5. Results

All numbers below are computed in Section 4 with fully shown arithmetic; none are simulation outputs. Values inherited from cited works are labeled as such.

Code parameters (exact, from equation (1)). For $p = 3$: $n_1 = 5$, $n_2 = 17$, $n_3 = 53$; for $p = 2$: $n_1 = 4$, $n_2 = 10$, $n_3 = 22$. The distance is $d = 2t + 1$ from equation (2), conditional on the branch-permutation symmetry of the encoder as stated in [12]. The flagship block $p = 3$, $t = 1$ has $n_1 = 5$ qudits, $d = 3$, and rate $1/5$.

Passive logical failure (exact computation, $p = 3$, $t = 1$, $\varepsilon = 10^{-3}$, binomial-tail convention). $P_{\text{pass}}(1) = 9.980 \times 10^{-6}$ per protection epoch, a suppression factor $S(1) = 100.2$ over an unencoded qudit. The leading-term truncation gives 9.970×10^{-6} ; the exact tail differs by 0.100%.

Passive logical failure ($p = 3$, $t = 2$, $\varepsilon = 10^{-3}$). $P_{\text{pass}}(2) = 6.729 \times 10^{-7}$ per protection epoch (leading term 6.705×10^{-7} plus a 0.35% correction from the weight-4 term), a suppression factor $S(2) = 1486$.

Decoder-model divergence (stated convention). Under the binomial-tail convention adopted in Section 4.3, the figures above hold. The alternative decoder model of [10] yields 5.992003×10^{-6} ($S = 166.9$) at $t = 1$ and 3.997×10^{-9} at $t = 2$; those values are not reproduced under our stated inputs and are reported here only as the divergent predecessor figures.

Thermal-activation projections (labeled projections; assumptions: barrier conjecture $E_{\text{barrier}}(d) = 2^d$ of [9], $k_B T = 1$, no tunnelling). $p_L(3) \leq 3.355 \times 10^{-4}$, $p_L(5) \leq 1.266 \times 10^{-14}$. Uncertainty: any tunnelling or coherent path across the 2^d -site cut invalidates these bounds entirely [9].

Thermodynamic ratio (computed under the assumption $T_h = 300$ K, ideal Carnot). The coefficient-of-performance ratio between 4 K and 10 mK operation is 405; removing a fixed heat load at 4 K requires 405 times less ideal input work than at 10 mK.

Threshold status. The 75% depolarizing threshold asserted in the ultrametric program [10] remains underived in this paper and in its predecessors; no computation here supports or refutes it. The classical iterated-majority evidence of [9] — suppression of $p = 0.4$ through 0.1736704, 0.0055485, and 2.905×10^{-10} over three levels — is consistent with a recursive threshold mechanism at the classical level but does not transfer to the quantum setting without Hypothesis Q.

6. Discussion

What is proved versus conjectured. Three results in this paper are exact: the ball-size identity (1), the binomial-tail failure probabilities at stated ε , and the Carnot ratio under the stated reservoir assumption. Two are conditional: the distance $d = 2t + 1$ holds only if a branch-permutation-symmetric encoder exists, and no explicit encoder is constructed here or in [12]; and the thermal-activation bounds hold only under the barrier conjecture of [9]. One central claim of the source program — the 75% depolarizing threshold — is neither derived nor refuted here; it remains an open conjecture, and the roughly 1% surface-code figure it is contrasted against is itself an assertion carried in the program documents [10, 11] rather than a result derived in this paper.

Limitations. The quantitative results rest on the following assumptions, with sensitivity noted for each:

- *Independent per-vertex depolarising noise* (Assumption A1). Correlated errors — a gate fault coupling two neighbouring vertices, or a burst affecting a whole branch — violate independence and could defeat the weight counting on which P_{pass} rests [12]. Sensitivity: a correlated two-vertex fault has the same effect as two independent faults, so modest correlations shift P_{pass} by factors of order the correlation strength; branch-wide bursts would destroy the suppression entirely.
- *Majority-vote decoder with symmetric encoder* (Assumption A2). If no symmetric encoder exists for a given ball, the weight- $\leq t$ correctability claim and the distance $d = 2t + 1$ fail [12]. Sensitivity: total — the failure probabilities computed here are meaningless without the encoder.
- *Noise level $\varepsilon = 10^{-3}$* . This is an assumption, chosen to match the predecessor analyses [9, 10]. Sensitivity: $P_{\text{pass}}(1) \approx 10 \binom{n_1}{2} \varepsilon^2$ scales quadratically, so at $\varepsilon = 10^{-2}$ the failure probability rises to roughly 9.9×10^{-4} and the suppression factor falls to about 10; the passive advantage over an unencoded qudit shrinks rapidly with ε .
- *Decoder-model convention*. The divergence between the 9.980×10^{-6} of [12] and the 5.992003×10^{-6} of [10] is unresolved in the literature; we take the conservative value, and all suppression factors in this paper are lower bounds relative to the [10] model.
- *Barrier conjecture and $k_B T = 1$* . The thermal-activation bounds are projections that any tunnelling or coherent path invalidates entirely [9].
- *Carnot ideality and $T_h = 300$ K*. Real cryocoolers operate at a small fraction of Carnot, with the fraction itself falling as T_c falls; the factor 405 is an upper bound on the real ratio.

Failure modes. Three ways the proposal can fail, following the failure-mode analysis of [9]: (i) the repetition-like quantum encoding does not correct X -type errors, so the classical majority-vote threshold does not transfer; (ii) coherent error paths interfere across branches, effectively reconverging paths that the ultrametric separates classically — this would falsify Hypothesis Q of [9]; (iii) the barrier is over-the-barrier only, and resonant tunnelling through the 2^d -site cut restores a depth-independent error channel. A fourth failure mode is specific to the encoder gap: the branch-permutation symmetry may be unachievable for any nontrivial ball, in which case the scheme reduces to a classical repetition structure with no quantum protection.

Falsification. The observation that would falsify the central quantitative result of this paper is a measured logical error rate above 2×10^{-5} per epoch for the five-qudit $p = 3$, $t = 1$ block under calibrated per-vertex depolarising noise at $\varepsilon = 10^{-3}$: our computed value 9.980×10^{-6} predicts a rate below that figure, and a reproducible excess would show that the binomial-tail model (Assumption A2) or the independence assumption (A1) fails in the physical implementation. Conversely, a measured rate at or below 10^{-5} would support the model at the stated precision. This criterion is the same one the analytic-derivation paper attaches to its five-atom proposal [10], and it is falsifiable with a single block and a calibrated noise source.

Open questions. The 75% threshold conjecture; the explicit construction of a symmetric encoder with verified Knill–Laflamme conditions $PE_i^\dagger E_j P = c_{ij} P$ for concrete Weyl error sets [12]; the quantum validity of Hypothesis Q [9]; and the reconciliation of the decoder-model divergence between [10] and [12] all remain open. No work cited here derives the 75% threshold or constructs the encoder, and we do not claim otherwise.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are quantum-hardware architects and fault-tolerance theorists evaluating whether a passive layer can replace or augment the syndrome-measurement layer of a stabilizer-code stack. If the result holds — a five-qudit block suppressing logical errors by a factor 100.2 at $\varepsilon = 10^{-3}$ with zero syndrome measurements — the design parameters that change are: the ancilla budget, which drops from the maximum stabilizer weight required by schemes such as Shor error correction [6] to zero; the classical decoding latency requirement, which drops from continuous real-time decoding to none; and the refrigeration specification, which relaxes from the 10–20 mK regime toward 4 K, worth up to a factor 405 in ideal cooling work per removed heat load (Section 4.6). Cloud quantum-service providers would see the effect as lower per-circuit cost if the passive block replaces active correction on a per-logical-qubit basis; this is a projection conditional on the encoder existing and on the threshold conjecture, not a derived economic figure.

(b) Current readiness. 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 formulated, with the ball-size identity and failure-probability computation established on paper but no experimental validation of any kind. The encoder construction, even in principle, has not been demonstrated, so no component or breadboard validation (TRL 3) is possible yet. No manufacturing readiness is claimed or claimable for theoretical work of this kind.

(c) Smallest next experiment. The minimal falsification test, following the five-atom proposal of [10], would be: a prototype of five ^{87}Rb atoms in an optical tweezer array arranged as a $p = 3$ tree of depth $t = 1$; each atom initialised in a chosen hyperfine level encoding a qudit; calibrated depolarising noise applied at $\varepsilon = 10^{-3}$; and a majority-vote readout of the root after one protection epoch. The measured logical error rate is the observable. A rate above 2×10^{-5} per epoch would falsify the binomial-tail model of Section 4.2; a rate at or below 10^{-5} would support it. Neutral-atom tweezer arrays are the natural platform because transversal operations and correlated decoding have already been demonstrated effective there [4], so the readout primitives exist. Before any hardware test, the smaller de-risking step is computational: an explicit symmetric encoder state vector for the five-vertex ball, with verification of $PE_i^\dagger E_j P = c_{ij} P$ for a concrete Weyl error set — a calculation that would confirm or kill the distance claim $d = 3$ at zero hardware cost.

(d) Binding constraints. In the order they would bind: first, the encoder — no symmetric encoder exists, and without one the scheme has no quantum content [12]; second, correlated noise — real two-qubit gates and crosstalk produce correlated faults that the independence assumption excludes, and the fault-injection methodology of [7] (inject faults, count escapes;

that work detected 100 of 100 message-fault injections and 100 of 100 application-fault injections with zero errors, reaching 97.6% summed coverage of protocol-deviation faults) is the template for characterizing whether tree-ball confinement survives correlations; third, qudit control — five-level hyperfine qudits require calibrated multi-level readout with error rates at or below the assumed $\varepsilon = 10^{-3}$, which is beyond typical current per-gate fidelities and is an open question rather than a derived requirement; fourth, Hypothesis Q — if coherent paths reconverge across branches, no amount of classical engineering recovers the suppression [9]. Cost drivers, scale-up economics, and regulatory constraints cannot be quantified from the inputs available and are left as open questions.

(e) Order-of-magnitude performance bounds (labeled projections). Under the assumptions of Section 4 ($\varepsilon = 10^{-3}$, binomial-tail decoder, symmetric encoder exists): a five-qudit block gives suppression 100.2; a seventeen-qudit block gives suppression 1486; scaling to $t = 3$ ($n_3 = 53$ qudits) would give, by the same quadratic-or-better tail scaling, a suppression factor of order 10^4 – 10^5 , though we have not computed the exact tail here. Under the additional barrier conjecture of [9] with $k_B T = 1$, depth-5 blocks would be bounded at $p_L \leq 1.266 \times 10^{-14}$. The refrigeration benefit is bounded above by the factor 405 (ideal Carnot, $T_h = 300$ K). All of these are projections conditional on the stated assumptions; none is a measurement.

(f) Risks and misuse paths. The principal risk is overclaiming by transfer: the classical majority-vote suppression of [9] and the classical Byzantine robustness of aggregation structures such as SignSGD [8] are real classical phenomena, and presenting them as evidence of quantum passive fault tolerance without an encoder and without Hypothesis Q would misrepresent the state of knowledge. A secondary risk is thermodynamic overclaim: the factor 405 is an ideal-Carnot bound, and quoting it as an achievable saving would overstate the case. The 75% threshold figure, if quoted without its conjectural status, would mislead system architects comparing against the surface-code baseline.

8. Conclusion

We formalized the ball-code model on the Bruhat-Tits tree T_{p+1} , derived the exact ball sizes $n_t = 1 + (p+1)(p^t - 1)/(p-1)$ giving $n_1 = 5$, $n_2 = 17$, $n_3 = 53$ for $p = 3$, and computed with fully shown arithmetic the passive logical failure probabilities $P_{\text{pass}}(1) = 9.980 \times 10^{-6}$ per epoch (suppression 100.2) for the five-qudit distance-3 block and $P_{\text{pass}}(2) = 6.729 \times 10^{-7}$ (suppression 1486) for the seventeen-qudit distance-5 block at $\varepsilon = 10^{-3}$, under the binomial-tail decoder convention chosen conservatively among the divergent predecessor models. Under the conjectured barrier $E_{\text{barrier}}(d) = 2^d$ of [9], thermal activation bounds the logical error rate at $p_L(3) \leq 3.355 \times 10^{-4}$ and $p_L(5) \leq 1.266 \times 10^{-14}$, both labeled projections invalidated by any tunnelling path. The Carnot coefficient-of-performance ratio between 4 K and 10 mK operation is 405, quantifying the cooling wall that passive operation would remove under fixed heat load and ideal refrigeration.

The honest summary is asymmetric: the combinatorial layer of the proposal is exact and cheap to verify, while the quantum layer rests on two unproved pillars — the existence of a branch-permutation-symmetric encoder and the confinement of superposed error paths (Hypothesis Q of [9]). The 75% threshold remains an open conjecture that neither this paper nor its predecessors derive. The path forward is ordered by cost: first an explicit encoder calculation for the five-vertex ball, then the five-atom tweezer experiment whose falsification threshold of 2×10^{-5} per epoch we state explicitly. If the encoder exists and the measured rate falls below that threshold, the passive scheme earns a place in the architectural comparison; until then, the contribution of this paper is a precise, conservative, and falsifiable quantification of what the tree geometry delivers on paper.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.39.2-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.

- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: 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] On the missing branches of the Bruhat-Tits tree. arXiv:1712.01463v2. <https://arxiv.org/abs/1712.01463v2> [2] Orthogonal Fault Tolerance for Dynamically Adaptive Systems. arXiv:1404.6830v1. <https://arxiv.org/abs/1404.6830v1> [3] Fault-tolerant gates via homological product codes. arXiv:1807.09783v2. <https://arxiv.org/abs/1807.09783v2> [4] Low-Overhead Transversal Fault Tolerance for Universal Quantum Computation. arXiv:2406.17653v2. <https://arxiv.org/abs/2406.17653v2> [5] Scheduling and Checkpointing optimization algorithm for Byzantine fault tolerance in Cloud Clusters. arXiv:1802.00951v1. <https://arxiv.org/abs/1802.00951v1> [6] Achieving fault tolerance on capped color codes with few ancillas. arXiv:2106.02649v4. <https://arxiv.org/abs/2106.02649v4> [7] Decentralized Validation for Non-malicious Arbitrary Fault Tolerance in Paxos. arXiv:2002.07509v1. <https://arxiv.org/abs/2002.07509v1> [8] SignSGD: Fault-Tolerance to Blind and Byzantine Adversaries. arXiv:2202.02085v2. <https://arxiv.org/abs/2202.02085v2> [9] <https://papers.qnfo.org/papers/ultrametric-quantum-error-correction-on-bruhattits-trees-energy-barriers-classical-evidence-and-the-quantization-question>. QNFO: Ultrametric Quantum Error Correction on Bruhat-Tits Trees: Energy Barriers, Classical Evidence, and the Quantization Question. [10] <https://papers.qnfo.org/papers/ultrametric-passive-fault-tolerance-on-bruhattits-trees-analytic-derivation-and-quantitative-projections>. QNFO: Ultrametric Passive Fault Tolerance on Bruhat-Tits Trees: Analytic Derivation and Quantitative Projections. [11] <https://papers.qnfo.org/papers/alpha-pi-project>. QNFO: Alpha Pi Project. [12] <https://papers.qnfo.org/papers/ultrametric-geometry-as-passive-quantum-error-correction-bruhattits-trees-qudit-codes>. QNFO: Ultrametric Geometry as Passive Quantum Error Correction: Bruhat-Tits Qudit Codes and the Non-Archimedean Bound.