

Scaling of the Hadamard Deformation Space in Concatenated $[[4,1,2]]$ Iceberg Codes: Exact Counts, Skew Amplification, and Search Reduction

Scaling of the Hadamard Deformation Space in Concatenated $[[4,1,2]]$ Iceberg Codes: Exact Counts, Skew Amplification, and Search Reduction

Abstract

Concatenated quantum error-correcting codes raise the effective distance of a joint code without requiring the discovery of new codes with desirable parameters, and the $[[4, 1, 2]]$ Iceberg code is the smallest error-detecting building block and therefore the one with the lowest resource overhead. The Iceberg code carries an asymmetric number of X and Z checks, so concatenating it produces skewed logical error rates between the X and Z observables; Clifford-deforming individual blocks can mitigate this, but the space of Hadamard deformations grows doubly exponentially with the number of concatenation levels, and no closed-form account of that growth is available in the grounding literature. The deformation space is counted under two equivalent conventions that are shown to coincide. The paper derives the exact deformation-space size $N_{\text{def}}(L) = 2^{(4^{L+1}-4)/3}$ for an L -level concatenation, giving 16, 1,048,576, and $\approx 1.93 \times 10^{25}$ candidate deformations at $L = 1, 2, 3$, and shows that a per-block greedy rule reduces the search at $L = 3$ from 2^{84} candidates to 336 evaluations, a reduction factor of $\approx 5.76 \times 10^{22}$. Under a stated phenomenological code-capacity model, unmitigated observable skew grows as η^{2^L} for physical bias η , and at $L = 2$, $\eta = 10$, $p_0 = 10^{-3}$ the projected total logical error rate is 2.160216×10^{-6} , more than 99.99% of which lies in the Z channel. No simulations were performed; all numbers are exact counts or labeled projections, and a minimal validation protocol is specified.

1. Introduction

Builders of fault-tolerant quantum memories face a concrete resource problem: they need codes whose logical error rates fall below the thresholds at which long computations accumulate uncorrectable failures, but codes combining large distance with low overhead are scarce. Concatenation sidesteps the scarcity by composing small known codes into a larger joint code whose effective distance grows with the number of levels, and the $[[4, 1, 2]]$ Iceberg code — the smallest error-detecting code — is the cheapest available building block for this construction [1]. A team selecting an inner code for a concatenated memory can adopt the Iceberg code at minimal qubit overhead, but the same choice imports a structural defect: the Iceberg code has an asymmetric number of X and Z checks, and concatenating it produces skewed logical error rates between the X and Z observables [1]. A memory whose two logical error channels fail at visibly different rates cannot be characterized by a single error budget; calibration, decoding, and logical-operation scheduling must all be handled per observable, and the worse channel dominates the memory lifetime.

The known remedy is to Clifford-deform the individual blocks of the concatenation — in the Hadamard-restricted case, to insert single-qubit Hadamard gates that swap X - and Z -type roles on selected physical qubits, rebalancing the check structure of the concatenated code. The obstacle is combinatorial: the space of possible Hadamard deformations grows doubly exponentially with the number of concatenation levels [1]. An exhaustive search that is feasible at one level is impossible at three, and without a closed-form count of the deformation space a designer cannot say in advance at which depth exhaustive search stops being viable. If the deformation space can be formalized and its growth derived explicitly, the decision that changes is the search policy: a designer can compute,

for each planned concatenation depth, exactly how many candidate deformations an exhaustive sweep would require, and can justify switching to a structured strategy-based or greedy search at a stated depth rather than discovering the intractability empirically.

The analysis is grounded in the source study of asymmetric QEC code concatenation [1]. The contributions are:

- A closed-form derivation of the Hadamard deformation-space size for an L -level concatenation of the $[[4, 1, 2]]$ Iceberg code, $N_{\text{def}}(L) = 2^{(4^{L+1}-4)/3}$, with explicit arithmetic at $L = 1, 2, 3, 4$ (Section 4.1), together with the demonstration that the per-internal-node model of one draft reproduces this count exactly when the per-node choice set has $c = 16$ elements (Section 4.1).
- A derivation of the physical-qubit overhead $n_L = 4^L$, the concatenation distance product rule $d_L = 2^L$, and the rate 4^{-L} , with worked values (Section 4.2).
- A phenomenological code-capacity model yielding the projected logical error rate $p_L = \alpha^{2^L-1} p_0^{2^L}$ and the skew amplification law η^{2^L} , with worked numbers under stated assumptions (Section 4.3).
- A greedy per-block search rule with an explicit evaluation count, reducing the $L = 3$ search from 2^{84} candidates to **336** evaluations under a stated factorization assumption (Section 4.4).
- A tractability frontier for exhaustive search as a function of a notional simulation budget (Section 4.4).
- A critical discussion of the assumptions the results rest on, of falsification criteria, and of the minimal next test (Sections 6 and 7).

2. Background and Related Work

The technical anchor of this paper is the source study on asymmetric QEC code concatenation [1]. Its supplied abstract establishes the facts this paper builds on: concatenated quantum error-correcting codes have recently gained popularity because they improve the effective distance of a joint code without requiring the discovery of new codes with desirable parameters; the $[[4, 1, 2]]$ Iceberg code is a common building block because it is the smallest error-detecting code and therefore has the lowest resource overhead; the Iceberg code has an asymmetric number of X and Z checks, and concatenating it produces skewed logical error rates between the X and Z observables; and the space of possible Hadamard deformations grows doubly exponentially with the number of concatenation levels. The study formalizes this deformation space, proposes six concrete deformation strategies, and compares them under a code capacity noise model with and without bias using a concatenated maximum-likelihood soft-information decoder; the abstract reports that one strategy achieves near-identical X and Z logical error rates in a logical-memory setting while reducing the total rate, but the supplied abstract is truncated mid-sentence at that clause, so no further detail about the identity of the winning strategy or its numerical margins is available from the grounding. This paper therefore derives the scaling structure around the study rather than reproducing its simulation results.

Entry [2] is the arXiv record for the same submission, identifier arXiv:2609.25251v1; its supplied summary repeats the opening of the same abstract and provides no additional technical detail beyond what is cited from [1]. It is retained here to document the retrieval provenance of the primary source.

Entries [3] through [8] are arXiv records whose supplied summaries state administrative withdrawal or removal and give no technical content. Entry [3], "A Simulation and Modeling of Access Points with Definition Language" (arXiv:1304.1836v2), was withdrawn by arXiv administrators because it contained fictitious content and was submitted under a pseudonym; its summary supplies no further detail. Entry [4], on superconductivity as a consequence of an ordering of the electron gas zero-point oscillations (arXiv:1005.0280v6), was administratively withdrawn as a duplicate of arXiv:1008.2691; the summary gives no technical detail. Entry [5], "Intuitionistic Fuzzy Ideals in Γ -semiring" (arXiv:1011.5746v2), was withdrawn due to plagiarized content from arXiv:1010.2469. Entry

[6], "Internal Location Based System For Mobile Devices Using Passive RFID And Wireless Technology" (arXiv:1001.2258v2), was withdrawn due to plagiarized content from arXiv:1009.3448. Entry [7], a comment titled "The meaning of systematic errors" on a paper concerning the Lense-Thirring effect detection with a Mars orbiter (arXiv:gr-qc/0703020v3), was removed because the submitter 'G. Felici' is an apparent pseudonym in violation of arXiv policies; its title indicates a topic in the metrology of systematic errors, which is thematically adjacent to the error-budgeting concerns of this paper, but the summary supplies no substantive content to build on. Entry [8], "Explicit estimates on prime numbers" (arXiv:1407.7158v2), was withdrawn due to fraudulent affiliation claims by the original submitter; its summary gives no further detail. These six entries are cited here to document the source-screening step of the methods (Section 3): none contributes a technical claim, and no statement in this paper depends on their content.

Entry [9] (arXiv:2006.16817v2) concerns a proof of monotonic increase in the cost function for the Krotov algorithm for open quantum systems. Its supplied summary states that a great number of quantum control papers have used one of the variants of the monotonically convergent variational control algorithm of Krotov, as described in works it cites from 2003, 1992, and 1998, and it references a 2018 paper on speeding up thermalisation via open quantum system variational optimisation; the summary is truncated and supplies no detail on the proof itself. The entry is relevant here only as an example of convergence-analysis work in the broader quantum control literature, and no claim of this paper rests on it.

The literature treatment above is deliberately constrained: of the nine supplied entries, only [1]/[2] carries substantive technical content on the paper's topic. Section 6 states this single-source anchoring as a limitation of the grounding.

3. Methods

The method is analytical combinatorics plus phenomenological modeling, applied to the concatenation structure described in the source study [1]. No simulations were run for this paper; all quantitative results are either exact combinatorial counts derived in Section 4 or clearly labeled projections of a stated noise model.

Source screening. The grounding block supplies nine bibliography entries. Each entry was screened on its own summary: entries whose summaries state administrative withdrawal, plagiarism, pseudonymity, or fictitious content ([3]–[8]) were excluded from technical use; the entry concerning another domain ([9]) was excluded from technical use; entry [1] (with duplicate record [2]) was retained as the sole technical anchor. In a corpus of this composition, single-source anchoring is not a choice but a necessity, and the provenance status of every non-anchor entry is recorded in Section 2.

Concatenation structure. The building block is the $[[4, 1, 2]]$ Iceberg code: $n_1 = 4$ physical qubits, $k = 1$ logical qubit, distance $d_1 = 2$; it detects errors but, with distance 2, does not correct an arbitrary single-qubit error. Concatenating L levels means composing the code with itself L times, so the physical-qubit count multiplies by 4 per level and the distance multiplies by 2 per level under the standard tensor-product rule for code concatenation, in which the distance of a concatenated code is the product of the distances of its factors. Both consequences are derived explicitly in Section 4.2.

Deformation model (adopted convention). A Hadamard deformation of a base block is a choice, for each of the block's $n_1 = 4$ physical qubits, of whether to apply a Hadamard gate H (which exchanges the roles of X - and Z -type checks on that qubit) or to leave the qubit undeformed. This gives $2^4 = 16$ deformations per base block. At level ℓ of the concatenation there are $4^{\ell-1}$ base blocks, each independently deformable, so the level- ℓ deformation count is $16^{4^{\ell-1}} = 2^{4^\ell}$, and a full deformation specification fixes a deformation for every base block at every level. Section 4.1 derives the resulting total count and verifies that it grows doubly exponentially, i.e., that

$\log_2 \log_2 N_{\text{def}}(L)$ is linear in L , consistent with the source study's characterization [1]. The adopted convention is the one in which every base block at every level is independently deformable, which is the literal reading of "deforming the individual blocks of the concatenation" [1]; a coarser convention with only L independent per-level choices would give a smaller count and is not adopted here.

Equivalence of the per-node convention. An equivalent tree model treats the concatenation as a tree with branching factor $b = 4$ and depth d , in which each internal node admits a choice among c distinct Clifford candidates, giving $c^{(b^d-1)/(b-1)}$ global assignments. With $c = 16$ this equals $16^{(4^d-1)/3} = 2^{4(4^d-1)/3} = 2^{(4^{d+1}-4)/3}$, which is exactly the adopted count with $L = d$: the two conventions are the same model, one counted per internal node with $c = 16$ candidates and one counted per base block with a binary choice per qubit. This identity is shown explicitly in Section 4.1.

Noise and decoding context. The source study compares deformation strategies under a code capacity noise model, with and without bias, using a concatenated maximum-likelihood soft-information decoder [1]. Code capacity noise means independent errors on physical qubits with rates p_X and p_Z for the two Pauli components; bias η denotes the ratio $\eta = p_Z/p_X$. This paper adopts the same noise-model family for its phenomenological projections (Section 4.3) but does not implement the decoder; the projections are analytic approximations whose assumptions are stated inline and whose role is to quantify the scaling of skew, not to predict absolute decoder performance.

Factorization assumption for the greedy rule. The greedy search rule of Section 4.4 rests on an explicit assumption, stated here as Assumption A:

Assumption A. To leading order in the small physical error probability, the total logical error rate of the concatenated code is a monotone increasing function of the sum over base blocks of the per-block error contributions, so that minimizing each block's contribution locally minimizes the total.

Assumption A is the standard leading-order factorization used in concatenation analyses: the lowest-weight logical operators of a concatenated code are supported on single low-weight paths through the concatenation tree, so the leading term of the logical error rate is a sum over paths of products of per-block errors, and the leading term of the total is monotone in the per-block sums. It is stated as an assumption because its validity range — which physical error rates, which depths — is not proven here and should be checked by the validation simulation of Section 7.

Definitions. The observable skew at level L is $\Delta_L = |p_X^{(L)} - p_Z^{(L)}|/(p_X^{(L)} + p_Z^{(L)})$, a dimensionless number in $[0, 1]$; $\Delta_L = 0$ is perfect symmetry and $\Delta_L = 1$ is complete one-sidedness. The total logical error rate is defined, following the source study, as the sum of the X and Z rates, $p_{\text{tot}}^{(L)} = p_X^{(L)} + p_Z^{(L)}$ [1].

4. Analysis

4.1 Deformation-space size

Input 1. The Iceberg block has $n_1 = 4$ physical qubits (from the code label $[[4, 1, 2]]$ [1]).

Input 2. A Hadamard deformation of a base block is a binary choice per qubit: deform or not. The source study states that the deformation space consists of Hadamard deformations of individual blocks and grows doubly exponentially with the number of concatenation levels [1]; the binary-per-qubit model is the minimal model consistent with that statement.

Step 1 (per-block count). A base block of 4 qubits admits $2^4 = 16$ Hadamard deformations.

Step 2 (blocks per level). At level ℓ the code occupies $n_\ell = 4^\ell$ physical qubits (derived in Section 4.2), which comprise $4^\ell/4 = 4^{\ell-1}$ base blocks. Each is independently deformable, so the level- ℓ deformation count is

$$N_\ell = 16^{4^{\ell-1}} = (2^4)^{4^{\ell-1}} = 2^{4^\ell}.$$

Step 3 (total count). A full deformation specification fixes a deformation at every base block of every level, so

$$N_{\text{def}}(L) = \prod_{\ell=1}^L 2^{4^\ell} = 2^{\sum_{\ell=1}^L 4^\ell} = 2^{\Sigma(L)}, \quad \Sigma(L) = \sum_{\ell=1}^L 4^\ell = \frac{4^{L+1} - 4}{3},$$

where the geometric sum is $\sum_{\ell=1}^L 4^\ell = 4(4^L - 1)/(4 - 1) = (4^{L+1} - 4)/3$.

Step 4 (worked values).

- $L = 1$: $\Sigma(1) = 4$, so $N_{\text{def}}(1) = 2^4 = 16$.
- $L = 2$: $\Sigma(2) = 4 + 16 = 20$, so $N_{\text{def}}(2) = 2^{20} = 1,048,576 \approx 1.05 \times 10^6$.
- $L = 3$: $\Sigma(3) = 4 + 16 + 64 = 84$, so $N_{\text{def}}(3) = 2^{84} = 19,342,813,113,834,066,795,298,816 \approx 1.93 \times 10^{25}$. (Check: $2^{80} = 1.2089258 \times 10^{24}$ and $2^{84} = 16 \times 2^{80} = 1.9342813 \times 10^{25}$.)
- $L = 4$: $\Sigma(4) = 84 + 256 = 340$, so $N_{\text{def}}(4) = 2^{340}$. In base 10, $340 \times \log_{10} 2 = 340 \times 0.301029995 = 102.350$, so $N_{\text{def}}(4) = 10^{102.350} \approx 2.24 \times 10^{102}$.

4.2 Concatenation overhead, distance, and rate

Input 3. The base block has $n_1 = 4$ physical qubits and $d_1 = 2$ (from the label $[[4, 1, 2]]$ [1]).

Step 1 (physical qubits). Each concatenation level replaces every physical qubit of the previous level by a full 4-qubit block, so the qubit count multiplies by 4 per level:

$$n_L = 4^L.$$

Worked values: $n_1 = 4, n_2 = 16, n_3 = 64, n_4 = 256$.

Step 2 (distance). Under the standard tensor-product rule for code concatenation, the distance of a concatenated code is the product of the distances of its factors (stated as the adopted convention in Section 3), so with $d_1 = 2$,

$$d_L = d_1^L = 2^L.$$

Worked values: $d_1 = 2, d_2 = 4, d_3 = 8, d_4 = 16$. The concatenated code detects up to $d_L - 1$ errors and corrects up to $\lfloor (d_L - 1)/2 \rfloor$ arbitrary errors.

Step 3 (rate). The concatenation encodes one logical qubit at every level ($k = 1$ throughout), so the encoding rate is

$$r_L = \frac{k}{n_L} = \frac{1}{4^L} = 4^{-L}.$$

Worked values: $r_1 = 0.25, r_2 = 0.0625, r_3 \approx 1.5625 \times 10^{-2}, r_4 \approx 3.90625 \times 10^{-3}$. The rate falls by a factor of 4 per level, which is the overhead cost of the distance gain 2 per level; this trade is what the source study's building block purchases [1].

4.3 Phenomenological skew amplification

Model inputs.

- **Input 4.** Code capacity noise with physical rates p_X and p_Z , bias $\eta = p_Z/p_X$ (Section 3, following the noise-model family of [1]).
- **Input 5.** The $[[4, 1, 2]]$ code has distance 2, so its weight-2 logical operators are supported on pairs of its 4 physical qubits; the number of such pairs is $\binom{4}{2} = 6$. This count is an exact combinatorial consequence of the code label [1] and the distance-2 structure.
- **Input 6 (model assumption, labeled).** The per-block logical error rate of one channel is αp^2 with $\alpha = 6$, i.e., the block fails in that channel when both qubits of one of the 6 weight-2 logical operators suffer errors of that Pauli type, and each such pair contributes independently. This is the standard leading-order code-capacity approximation for a distance-2 block; it neglects correlations and higher-weight contributions and is therefore a projection, not a decoder prediction.

Step 1 (per-level recursion). At each concatenation level, every logical qubit of the previous level is replaced by a block whose logical failure requires two of its constituent (now logical) failures on a weight-2 logical operator. With 6 such operators per block, the recursion for one channel with level- $\ell - 1$ rate $p^{(\ell-1)}$ is

$$p^{(\ell)} = 6 \left(p^{(\ell-1)} \right)^2.$$

Step 2 (closed form). Unrolling the recursion from the physical rate p_0 down to level L : the exponent of p_0 doubles plus one at each step, giving 2^L , and the factor 6 appears once per unrolling step, giving 6^{2^L-1} :

$$p^{(L)} = \alpha^{2^L-1} p_0^{2^L}, \quad \alpha = 6.$$

Verification of the unrolling at $L = 2$: $p^{(1)} = 6p_0^2$; $p^{(2)} = 6(6p_0^2)^2 = 6 \cdot 36 p_0^4 = 6^3 p_0^4 = \alpha^{2^2-1} p_0^{2^2}$. Consistent.

Step 3 (skew law). Applying the closed form to each channel separately,

$$\frac{p_Z^{(L)}}{p_X^{(L)}} = \frac{\alpha^{2^L-1} p_Z^{2^L}}{\alpha^{2^L-1} p_X^{2^L}} = \left(\frac{p_Z}{p_X} \right)^{2^L} = \eta^{2^L}.$$

The common prefactor α^{2^L-1} cancels, so the skew ratio amplifies as η^{2^L} independently of α : the asymmetry of the building block is compounded once per effective doubling, which is the analytic form of the statement that concatenating the Iceberg code produces skewed logical error rates [1].

Step 4 (worked numbers at $\eta = 10, p_0 = 10^{-3}$). Then $p_X = 10^{-3}$ and $p_Z = \eta p_X = 10^{-2}$.

- $L = 1$: $p_X^{(1)} = 6 \times (10^{-3})^2 = 6 \times 10^{-6}$; $p_Z^{(1)} = 6 \times (10^{-2})^2 = 6 \times 10^{-4}$. Total $p_{\text{tot}}^{(1)} = 6.06 \times 10^{-4}$. Skew ratio $\eta^2 = 10^2 = 100$.
- $L = 2$: $p_X^{(2)} = 6^3 \times (10^{-3})^4 = 216 \times 10^{-12} = 2.16 \times 10^{-10}$; $p_Z^{(2)} = 216 \times (10^{-2})^4 = 216 \times 10^{-8} = 2.16 \times 10^{-6}$. Total:

$$p_{\text{tot}}^{(2)} = 2.16 \times 10^{-6} + 2.16 \times 10^{-10} = 2.160216 \times 10^{-6}.$$

The Z share of the total is $2.16 \times 10^{-6} / 2.160216 \times 10^{-6} = 0.999900 \dots$, i.e., more than 99.99% of the total logical error rate lies in the Z channel. Skew ratio $\eta^4 = 10^4$.

- $L = 3$: $6^7 = 279936$. $p_X^{(3)} = 279936 \times (10^{-3})^8 = 2.79936 \times 10^{-19}$; $p_Z^{(3)} = 279936 \times (10^{-2})^8 = 2.79936 \times 10^{-11}$. Total $p_{\text{tot}}^{(3)} \approx 2.7993600 \times 10^{-11}$ (the X contribution adds 2.79936×10^{-19} , i.e., one part in 10^8). Skew ratio $\eta^8 = 10^8$.

Step 5 (observable skew). With $\Delta_L = |p_X^{(L)} - p_Z^{(L)}| / (p_X^{(L)} + p_Z^{(L)})$: at $L = 2$, $\Delta_2 = (2.16 \times 10^{-6} - 2.16 \times 10^{-10}) / 2.160216 \times 10^{-6} = 0.9998 \dots$; at $L = 3$, $\Delta_3 = 1 - 10^{-8}$. The skew observable approaches complete one-sidedness within three levels at bias $\eta = 10$ under this model, which quantifies why deformation is needed rather than optional.

4.4 Greedy search rule and tractability frontier

Input 7. Assumption A (Section 3): to leading order, the total logical error rate is monotone increasing in the sum over base blocks of the per-block error contributions, so per-block local minimization is a valid greedy objective.

Step 1 (greedy evaluation count). Under Assumption A, the search decomposes: each base block is optimized independently by sweeping its own 16 Hadamard deformations and keeping the one that minimizes that block's contribution. The number of base blocks across all levels of an L -level concatenation is

$$B(L) = \sum_{\ell=1}^L 4^{\ell-1} = \frac{4^L - 1}{3}.$$

At $L = 3$: $B(3) = (64 - 1)/3 = 21$ blocks ($1 + 4 + 16$). Each block requires 16 evaluations, so the greedy search costs

$$G(L) = 16 B(L) = 16 \times \frac{4^L - 1}{3}.$$

At $L = 3$: $G(3) = 16 \times 21 = 336$ evaluations.

Step 2 (reduction factor). The exhaustive count at $L = 3$ is $N_{\text{def}}(3) = 2^{84}$ (Section 4.1). The reduction factor is

$$\frac{N_{\text{def}}(3)}{G(3)} = \frac{2^{84}}{336} = \frac{2^{84}}{2^4 \times 21} = \frac{2^{80}}{21} = \frac{1208925819614629174706176}{21} \approx 5.7568 \times 10^{22}.$$

Arithmetic check: $21 \times 5.7568 \times 10^{22} = 1.20893 \times 10^{24} = 2^{80}$, correct to the digits shown. The greedy rule therefore reduces the $L = 3$ search from $2^{84} \approx 1.93 \times 10^{25}$ candidates to 336 evaluations, a reduction factor of $\approx 5.76 \times 10^{22}$.

Step 3 (tractability frontier). Let the notional evaluation budget be R evaluations per second and the tolerable sweep time T seconds. Exhaustive search is viable iff $N_{\text{def}}(L) \leq RT$. Taking $R = 10^6$ evaluations per second as a notional throughput for a cheap per-candidate evaluation (a labeled projection, not a measured figure):

- $L = 1$: $N_{\text{def}} = 16$; sweep time 1.6×10^{-5} s. Trivially viable.
- $L = 2$: $N_{\text{def}} = 2^{20} = 1,048,576$; sweep time ≈ 1.05 s. Viable.
- $L = 3$: $N_{\text{def}} = 2^{84} \approx 1.93 \times 10^{25}$; sweep time $\approx 1.93 \times 10^{19}$ s $\approx 6.13 \times 10^{11}$ years (using 1 year $\approx 3.156 \times 10^7$ s). Not viable.
- $L = 4$: $N_{\text{def}} = 2^{340} = 10^{102.350}$; sweep time $\approx 10^{96.35}$ s $\approx 10^{88.9}$ years. Not viable by an overwhelming margin.

The frontier therefore sits between $L = 2$ and $L = 3$ at this notional throughput, and it moves by at most a few orders of magnitude for any plausible R : since $N_{\text{def}}(3)/N_{\text{def}}(2) = 2^{64} \approx 1.8 \times 10^{19}$, no realistic throughput increase closes the $L = 3$ gap. The decision rule that follows is: exhaustive sweep at $L \leq 2$, greedy per-block

search from $L = 3$ upward.

5. Results

All results below are either exact combinatorial counts (derived in Sections 4.1, 4.2, 4.4) or projections of the stated code-capacity model with $\alpha = 6$ (derived in Section 4.3); no simulations were performed.

R1 (deformation-space size). The Hadamard deformation space of an L -level $[[4, 1, 2]]$ concatenation has exact size

$$N_{\text{def}}(L) = 2^{(4^{L+1}-4)/3},$$

with $N_{\text{def}}(1) = 16$, $N_{\text{def}}(2) = 1,048,576$, $N_{\text{def}}(3) = 2^{84} \approx 1.93 \times 10^{25}$, and $N_{\text{def}}(4) = 2^{340} \approx 2.24 \times 10^{102}$. The growth is doubly exponential: $\log_2 \log_2 N_{\text{def}}(L) = \log_2((4^{L+1}-4)/3)$ is asymptotically linear in L , consistent with the source study's characterization [1].

R2 (convention equivalence). The per-internal-node tree model with branching factor $b = 4$ and $c = 16$ candidates per node yields $c^{(b^d-1)/(b-1)} = 2^{(4^{d+1}-4)/3}$, identical to $N_{\text{def}}(L)$ at $L = d$ (Section 4.1 and Section 3); the two counting conventions are one model.

R3 (overhead). $n_L = 4^L$ physical qubits, $d_L = 2^L$ distance, $r_L = 4^{-L}$ rate; e.g., $L = 3$ gives 64 qubits, distance 8, rate $\approx 1.5625 \times 10^{-2}$.

R4 (skew amplification). Under the stated model, the observable skew ratio amplifies as η^{2^L} : at $\eta = 10$, the Z -to- X logical error ratio is $10^2, 10^4, 10^8$ at $L = 1, 2, 3$. At $L = 2$ with $p_0 = 10^{-3}$, the projected total logical error rate is 2.160216×10^{-6} , of which more than 99.99% is in the Z channel; at $L = 3$ the projected total is $\approx 2.80 \times 10^{-11}$ with skew ratio 10^8 . These are projections under Assumption A's model family and the $\alpha = 6$ leading-order approximation, not decoder measurements.

R5 (search reduction). The greedy per-block rule costs $G(L) = 16(4^L - 1)/3$ evaluations: 16, 80, 336 at $L = 1, 2, 3$. At $L = 3$ this reduces the search from 2^{84} candidates to 336 evaluations, a reduction factor of $2^{80}/21 \approx 5.76 \times 10^{22}$, valid to the extent Assumption A holds.

R6 (tractability frontier). At a notional throughput of 10^6 evaluations per second, exhaustive search is viable at $L \leq 2$ (≈ 1.05 s at $L = 2$) and non-viable at $L \geq 3$ ($\approx 6.13 \times 10^{11}$ years at $L = 3$); the gap of a factor $2^{64} \approx 1.8 \times 10^{19}$ between $L = 2$ and $L = 3$ cannot be closed by throughput improvements alone.

6. Discussion

The central result — the exact count $N_{\text{def}}(L) = 2^{(4^{L+1}-4)/3}$ and the resulting $L = 2 \rightarrow 3$ tractability frontier — is a combinatorial statement and is robust; the phenomenological results are not, and this section separates the two.

What is secure. The deformation count follows from three inputs only: the block size $n_1 = 4$, the binary-per-qubit deformation model, and independent deformability of every base block at every level. The first is fixed by the code label [1]; the second is the minimal model consistent with the source study's statement that the deformation space consists of Hadamard deformations of individual blocks and grows doubly exponentially [1]; the third is the literal reading of "deforming the individual blocks of the concatenation" [1], as the literal reading of the source study's statement requires. If a different deformation model were intended — for example, deformations restricted to a

subset of qubits, or strategies that deform blocks jointly — the count would change, but any model in which each of $4^{\ell-1}$ blocks at level ℓ carries an independent binary-per-qubit choice gives the same count. The equivalence result R2 shows the count does not depend on the choice between the two counting conventions.

What is provisional. The skew law η^{2^L} and the projected rates rest on the $\alpha = 6$ leading-order approximation and Assumption A. The approximation counts only weight-2 logical-operator failure paths and assumes independence of the two physical errors on each path; it neglects error correlations, degenerate failure structure of the Iceberg code, and the effect of the actual decoder. The source study uses a concatenated maximum-likelihood soft-information decoder under code capacity noise with and without bias [1]; a maximum-likelihood decoder can exploit correlations that the leading-order product form ignores, in which case the true rates would deviate from the projections — though the qualitative claim that the Z channel dominates at bias $\eta = 10$ follows from the raw ratio η^{2^L} alone and does not depend on α .

Falsification. The observation that would falsify the main quantitative projection is a decoder-level measurement at $L = 2$, $\eta = 10$, $p_0 = 10^{-3}$ in which the Z share of the total logical error rate falls substantially below the projected 99.99%, or at which the total rate departs from $\alpha^3 p_0^4$ by more than the uncertainty introduced by higher-weight paths. The observation that would falsify the structural claim R1 is a deformation-space formalization in the source study's full treatment that counts fewer or more choices per block than the binary-per-qubit model — for example, if the six deformation strategies of [1] implicitly restrict the per-block choice set to fewer than 16 options, the exhaustive counts would shrink accordingly, though the doubly exponential growth and the $L = 2 \rightarrow 3$ frontier would survive any fixed finite choice set.

Limitations list.

- *Binary-per-qubit deformation model.* The count N_{def} scales as $c^{(4^L-1)/3}$ in a general choice-set model with c options per block; the paper's numbers use $c = 16$. Sensitivity: linear in the exponent, so halving c to 8 changes $N_{\text{def}}(3)$ from 2^{84} to 2^{63} — still doubly exponential and still beyond exhaustive reach.
- *Independent deformability of all blocks at all levels.* If the source study's strategies constrain blocks jointly, the effective count is smaller. Sensitivity: the frontier location shifts, but not the doubly exponential character.
- *Assumption A (factorization).* The greedy count of 336 evaluations is valid only if per-block local minimization tracks the global optimum. Sensitivity: if Assumption A fails, the greedy result is a heuristic with no guarantee, and the honest fallback is the exhaustive sweep at $L \leq 2$ plus the source study's strategy-based search [1] at $L \geq 3$.
- *Leading-order rate model ($\alpha = 6$).* Neglects higher-weight paths and correlations. Sensitivity: absolute projected rates scale as α^{2^L-1} , so a factor-2 error in α becomes a factor $2^7 = 128$ at $L = 3$; the skew ratio η^{2^L} is insensitive to α because it cancels.
- *Single-source grounding.* Of the supplied bibliography, only [1]/[2] carries technical content on this topic; entries [3]–[9] are withdrawn, removed, or out-of-domain records with no usable technical content, and entries [10] and [13] concern a different research program. Every external claim therefore traces to one abstract, which is truncated mid-sentence and reports no numerical margins.
- *No simulations.* All rate numbers are projections; none has been checked against a decoder run.

Open questions. Whether the six deformation strategies of [1] correspond to subsets of the 16-element per-block choice set; whether Assumption A holds at the physical error rates where concatenated Iceberg codes are operated; and how the skew law interacts with a biased-noise decoder that reweights X and Z differently.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The direct user is a designer of concatenated fault-tolerant quantum memories who is selecting an inner code and a deformation policy: the source study's setting is exactly a logical-memory experiment built from $[[4, 1, 2]]$ blocks [1]. If the results hold, two decisions change. First, the search-policy decision: the designer can fix exhaustive deformation search for concatenations of depth $L \leq 2$ (at most 1,048,576 candidates, sweepable in seconds at any plausible throughput) and commit to a structured strategy-based or greedy search from $L = 3$ upward, on the basis of the exact count 2^{84} rather than an empirical discovery of intractability. Second, the error-budgeting decision: the skew law η^{2^L} tells the designer that at bias $\eta = 10$ the Z channel carries more than 99.99% of the projected logical error budget at $L = 2$, so calibration and decoding effort should be allocated per observable, with the Z channel treated as the lifetime-limiting channel, unless deformation restores symmetry as the source study reports for one of its strategies [1].

(b) Readiness. This work is purely analytical and has not been peer reviewed: by the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and reported; concept formulated). The exact counts are mathematics and do not mature further without implementation; no hardware, decoder software, or manufacturing readiness is claimed, and analytical work of this kind does not exceed TRL 3 even after validation against simulation.

(c) Smallest next test. The minimal de-risking experiment is a decoder-level simulation at $L = 2$ under code capacity noise with bias $\eta = 10$, $p_0 = 10^{-3}$: run a concatenated maximum-likelihood soft-information decoder of the kind the source study describes [1] over the full $2^{20} = 1,048,576$ deformation space (viable per R6) and measure the total logical error rate and its Z share. The outcome that refutes the projection is a measured Z share below, say, 99% of the total, or a measured total rate inconsistent with the $\alpha^3 p_0^4$ scaling beyond the spread induced by higher-weight paths. The same run tests Assumption A at zero additional cost: compare the greedy-selected deformation (336 per-block evaluations) against the exhaustive optimum; if the greedy choice's total rate exceeds the exhaustive optimum's by more than statistical uncertainty at realistic sample sizes, Assumption A fails at $L = 2$ and the greedy rule's use at $L \geq 3$ is unsupported.

(d) Binding constraints. In order of first binding: (i) compute for the validation simulation — a maximum-likelihood decoder over 2^{20} deformations at adequate sample size is a real but modest compute commitment, and it is the first bottleneck encountered; (ii) the validity range of Assumption A, which is unknown and bounds how far the greedy rule can be trusted beyond the tested depth; (iii) decoder fidelity — the projections use code capacity noise, and hardware noise with correlations, leakage, or measurement error would require re-derivation of the skew law's prefactors; (iv) whether the source study's six strategies [1] span or restrict the 16-element per-block choice set, which determines whether strategy-based search and per-block greedy search explore the same space. Materials, tolerances, and component constraints do not bind at this stage because no hardware is proposed.

(e) Order-of-magnitude bounds (projections). With the notional throughput $R = 10^6$ evaluations per second stated in Section 4.4: exhaustive sweep cost is ≈ 1.05 s at $L = 2$ and $\approx 6.13 \times 10^{11}$ years at $L = 3$; the greedy cost is 336 evaluations, i.e., $\approx 3.4 \times 10^{-4}$ s at the same notional rate. These figures assume per-candidate evaluation cost of 10^{-6} s, which is an assumption about a cheap per-block objective, not a measured decoder runtime; a full maximum-likelihood evaluation per candidate would raise all times by the ratio of the true per-candidate cost to 10^{-6} s and could move the $L = 2$ sweep from seconds to hours without changing the frontier's location between $L = 2$ and $L = 3$, since the $L = 3$ gap is a factor $\approx 1.8 \times 10^{19}$.

(f) Risks and misuse paths. The main misuse is reading the projected rates $p_{\text{tot}}^{(L)}$ of Section 4.3 as predictions of a real decoder's performance. They are leading-order code-capacity projections with the $\alpha = 6$ approximation and Assumption A built in, and a design decision — for example, skipping X -channel calibration because the model

says the X share is below 10^{-4} of the total at $L = 2$ — would be unjustified until the validation run of part (c) confirms the projected shares. A second misuse is treating the greedy rule as an optimizer with a guarantee: its 336-evaluation cost at $L = 3$ is valid only under Assumption A, and applying it at depths where Assumption A has not been tested could silently return a deformation whose total logical error rate is far from the global optimum, which is worse than an honest exhaustive sweep at a shallower depth. A third misuse is extrapolating the tractability frontier beyond its stated throughput assumption: the frontier location between $L = 2$ and $L = 3$ is robust to throughput because the gap is a factor $2^{64} \approx 1.8 \times 10^{19}$, but the absolute sweep times in part (e) scale linearly with the assumed per-candidate cost and should not be quoted without restating that assumption.

8. Conclusion

This paper derived the exact size of the Hadamard deformation space for an L -level concatenation of the $[[4, 1, 2]]$ Ice code: $N_{\text{def}}(L) = 2^{(4^{L+1}-4)/3}$, giving 16, 1,048,576, and $\approx 1.93 \times 10^{25}$ candidate deformations at $L = 1, 2, 3$. The count follows from three inputs only — the block size $n_1 = 4$ fixed by the code label [1], a binary-per-qubit deformation model, and independent deformability of every base block at every level — and it was shown to coincide exactly with a per-internal-node tree model with 16 candidates per node, so the result does not depend on the choice between the two counting conventions. The count makes the source study's statement that the deformation space grows doubly exponentially with the number of concatenation levels [1] quantitative, and it locates the exhaustive-search frontier: at any plausible evaluation throughput, an exhaustive sweep is viable at $L \leq 2$ and impossible at $L \geq 3$, because the step from $L = 2$ to $L = 3$ multiplies the search space by $2^{64} \approx 1.8 \times 10^{19}$.

On top of the count, the paper derived the overhead structure of the concatenation ($n_L = 4^L$, $d_L = 2^L$, $r_L = 4^{-L}$), a phenomenological skew amplification law under which the ratio of Z to X logical error rates grows as η^{2^L} with physical bias η , and a greedy per-block search rule that reduces the $L = 3$ search from 2^{84} candidates to 336 evaluations under an explicitly stated factorization assumption. At $L = 2$, $\eta = 10$, $p_0 = 10^{-3}$, the projected total logical error rate is 2.160216×10^{-6} , of which more than 99.99% lies in the Z channel — a quantification of why the Iceberg code's asymmetric check structure [1] makes deformation a necessity rather than an option at depth.

The secure part of the paper is the combinatorics; the provisional part is everything that depends on the $\alpha = 6$ leading-order rate model and Assumption A. The single decisive test is the one specified in Section 7(c): a decoder-level simulation at $L = 2$ over the full 2^{20} deformation space, measuring the total logical error rate and its Z share, which simultaneously checks the skew projection and the greedy rule's validity at zero additional cost. If that test confirms the projections, the practical yield is a stated search policy — exhaustive below $L = 3$, structured or greedy from $L = 3$ upward — grounded in an exact count rather than an empirical discovery of intractability; if it refutes them, the count $N_{\text{def}}(L)$ survives regardless, since it rests on no rate model at all.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.16.0-revise-retry). 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; statements about a cited work are restricted to its retrieved title and abstract.
- 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 17 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 its abstract; 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] TITLE: arXiv Query: `searchquery=&idlist=2609.25251&start=0&max_results=1` [2] Exploring Asymmetric QEC Code Concatenation. arXiv:2609.25251v1. <https://arxiv.org/abs/2609.25251v1> [3] A Simulation and Modeling of Access Points with Definition Language. arXiv:1304.1836v2. <https://arxiv.org/abs/1304.1836v2> [4] Superconductivity as a consequence of an ordering of the electron gas zero-point oscillations. arXiv:1005.0280v6. <https://arxiv.org/abs/1005.0280v6> [5] Intuitionistic Fuzzy Ideals in Γ -semiring. arXiv:1011.5746v2. <https://arxiv.org/abs/1011.5746v2> [6] Internal Location Based System For Mobile Devices Using Passive RFID And Wireless Technology. arXiv:1001.2258v2. <https://arxiv.org/abs/1001.2258v2> [7] The meaning of systematic errors, a comment to "Reply to On the Systematic Errors in the Detection of the Lense-Thirring Effect with a Mars Orbiter", by Lorenzo Iorio. arXiv:gr-qc/0703020v3. <https://arxiv.org/abs/gr-qc/0703020v3> [8] Explicit estimates on prime numbers. arXiv:1407.7158v2. <https://arxiv.org/abs/1407.7158v2> [9] Proof of monotonic increase in the cost function for Krotov algorithm for open quantum systems. arXiv:2006.16817v2. <https://arxiv.org/abs/2006.16817v2> [10] QNFO: QNFO Consilient Synthesis v2.0: The Shared Adelic Kernel [13] QNFO: The Consilience Framework: From Valuation Theory to the Void — A Cross-Domain Synthesis