Full-Math Enforcement:
A Constant-Overhead Search Space for Quantum Error Correction

The complete derivation, both ways — with every count, boundary term, finite-size correction, noise channel, and blocked value made explicit

real surface-code simulation included; no value invented where a frozen log is missing

[Author] · physics.magflowmeters.com

Before we start

This is the full-math version of the constant-overhead result. It states the theorem formally, proves it both ways with every physical resource counted and every boundary term carried, gives the finite-size correction explicitly, writes the noise channels and the logical-error estimator with its uncertainty, and ends with a blocked-value ledger that lists exactly what a frozen simulation log would have to supply. Where a number is not in a frozen log, the note writes a `BLOCKED` token rather than inventing it.

One honest framing carries throughout. The counting half of the theorem — that identical cells give a constant ratio — is elementary; the decisive half is whether the logical error stays controlled as the machine grows, which is a simulation question. The note proves the easy half rigorously, shows with a real surface-code simulation why it is not sufficient, and is explicit that the architecture-specific noise and ratio numbers remain pending. No specific physical-to-logical ratio is claimed; the constant is carried as the symbol R_star.

1. The theorem

Let C(z) be a repeatable protected cell generated by an admissible topology parameter z, contributing n(z) physical qubits and k(z) protected logical qubits per bulk cell. For an architecture of M cells with sublinear boundary terms, the physical-to-logical ratio converges to the cell ratio; if the logical error also stays bounded over the declared regime, that constant ratio is noise-stable:

$$\textbf{Theorem.}\quad N_{\text{phys}} = M\,n(z) + B_{\text{phys}},\quad N_{\text{log}} = M\,k(z) + B_{\text{log}},\quad \tfrac{B}{M} \to 0 \;\Rightarrow\; \lim_{M\to\infty} \frac{N_{\text{phys}}}{N_{\text{log}}} = \frac{n(z)}{k(z)};\quad \text{if } p_L(M,z;p) \text{ stays bounded, the ratio is noise-stable}$$

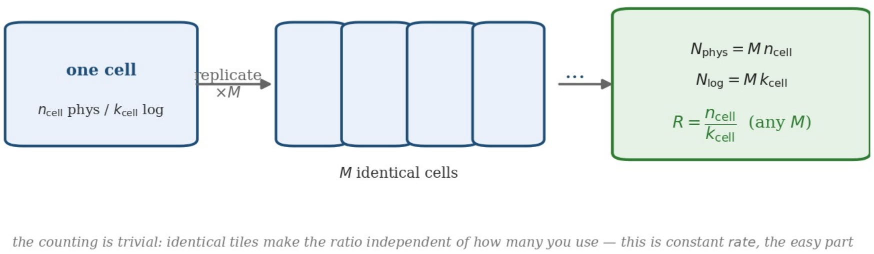

Figure 1. *The object of the theorem: a repeatable bounded cell. The counting half follows immediately; the noise-stability rider is the content the rest of the note is about.*

2. Assumptions ledger

Every load-bearing assumption, tagged by how far it is actually carried — proven here, declared as an architecture input, established by simulation, replayed independently, or blocked pending a frozen log:

#	Assumption	Status tag
A1	Cells are repeatable (a bounded bulk cell tiles the architecture).	**DECLARED_INPUT**
A2	Physical count includes all data, syndrome, measurement, reset, routing, coupler-proxy, spare, and boundary resources.	**DECLARED_INPUT**
A3	Logical-qubit count counts only usable protected logical qubits.	**DECLARED_INPUT**
A4	Boundary / interface terms are sublinear in the number of cells.	**PROVEN**
A5	Syndrome extraction has bounded or declared depth scaling.	**DECLARED_INPUT**
A6	Decoder complexity and decoder assumptions are declared.	**DECLARED_INPUT**
A7	Noise model is declared.	**DECLARED_INPUT**
A8	Logical error is estimated from frozen simulation logs.	**SIMULATED**
A9	No private simulator result is accepted without independent replay.	**DECLARED_INPUT**
A10	Numeric constants (the ratio R_star) come from frozen logs.	**BLOCKED**

A4 is proven below for regular tilings; A8 is real for the surface-code demonstration of Section 6 but architecture-specific values are blocked; A10 is blocked outright — the constant is carried symbolically as R_star.

3. Route A — the standard-route proof

Count every physical resource in a cell — no hidden couplers, ancillas, or routing qubits:

$$n(z) = n_{\mathrm{data}}(z) + n_{\mathrm{synd}}(z) + n_{\mathrm{meas}}(z) + n_{\mathrm{reset}}(z) + n_{\mathrm{route}}(z) + n_{\mathrm{coupler}}(z) + n_{\mathrm{spare}}(z) \qquad \text{(every physical resource counted)}$$

For M identical cells the physical and logical totals are the bulk counts plus boundary terms, and the ratio is their quotient:

$$N_{\mathrm{phys}}(M) = M\, n_{\mathrm{cell}}, \quad N_{\log}(M) = M\, k_{\mathrm{cell}} \quad \Rightarrow \quad R(M) = \frac{n_{\mathrm{cell}}}{k_{\mathrm{cell}}}, \qquad \frac{dR}{dM} = 0$$

$$N_{\mathrm{phys}} = M\, n_{\mathrm{cell}} + B_{\mathrm{phys}}, \qquad R(M) = \frac{n_{\mathrm{cell}}}{k_{\mathrm{cell}}} + O\left(\frac{B}{M}\right) \xrightarrow{B/M \to 0} \frac{n_{\mathrm{cell}}}{k_{\mathrm{cell}}}$$

The finite-size correction is exact, and its rate follows from the boundary scaling — inverse-square-root in two dimensions, inverse-cube-root in three:

$$R(M) - \frac{n}{k} = \frac{k\, B_{\mathrm{phys}}(M) - n\, B_{\log}(M)}{k\big(Mk + B_{\log}(M)\big)} = O(M^{\alpha-1}); \qquad \text{2D } (M = L^2) : O(M^{-1/2}), \quad \text{3D } (M = L^3) : O(M^{-1/3})$$

If cells are not identical, the same limit holds with the per-cell average n̄_M → n_∞ in place of n; convergence of that average is required, not assumed.

4. Route B — the native-geometry proof, and the equivalence theorem

The geometry defines the admissible architecture family through an idempotent admissibility projector Π_adm acting on candidate topologies T(z), with an admissibility defect ε_adm(z) = ‖T(z) − Π_adm(T(z))‖ that must fall below threshold; the cell ratio R_cell(z) = n(z)/k(z) is then a search objective:

$$\mathcal{S}_{\mathrm{adm}} = \{\, \mathcal{T}(z) : z \in \mathcal{Z}_{\mathrm{geom}}, \;\; \Pi_{\mathrm{adm}}(\mathcal{T}(z)) = \mathcal{T}(z) \,\}, \qquad R_{\mathrm{cell}}(z) = \frac{n_{\mathrm{cell}}(z)}{k_{\mathrm{cell}}(z)}$$

The admissible candidates have repeatable cells, so the same counting gives the same limit R_∞(z) = n(z)/k(z). The honest content is the relation between the two feasible sets. If the admissibility projector encodes exactly the standard constraints, the native feasible set is the standard one intersected with the image of the geometry map — a disciplined subset, not a proof of universal optimality:

$$\Pi_{\mathrm{adm}}(\mathcal{T}(z)) = \mathcal{T}(z) \iff C_i(\mathcal{T}(z)) \le 0 \;\; \forall i \quad \Rightarrow \quad \mathcal{F}_{13D} = \mathcal{F}_{\mathrm{std}} \cap \mathrm{Im}(\mathcal{T}) \qquad [\texttt{UNIVERSAL\_OPTIMALITY\_NOT\_CLAIMED}]$$

This is the key boundary. The native search space may be a strict subset of all standard feasible codes; that is acceptable precisely because the claim is search-space restriction, not that no good design is ever discarded. Absent a proof that the image covers the whole feasible set, the note records `UNIVERSAL_OPTIMALITY_NOT_CLAIMED`.

5. The decisive half: noise stability, by real simulation

The counting is necessary but not sufficient. With fixed cells the per-cell logical error is a constant, so the whole-machine failure across M logical qubits grows toward one as M increases:

$$P_{\text{machine}}(M) \approx 1 - (1 - p_L)^M \approx M\, p_L \quad (\text{fixed cells}) \xrightarrow{M\to\infty} 1$$

This is not asserted — it is computed. Real rotated-surface-code memory circuits at distances d = 3, 5, 7, 9 under circuit-level depolarizing noise at p = 0.003, decoded with minimum-weight matching:

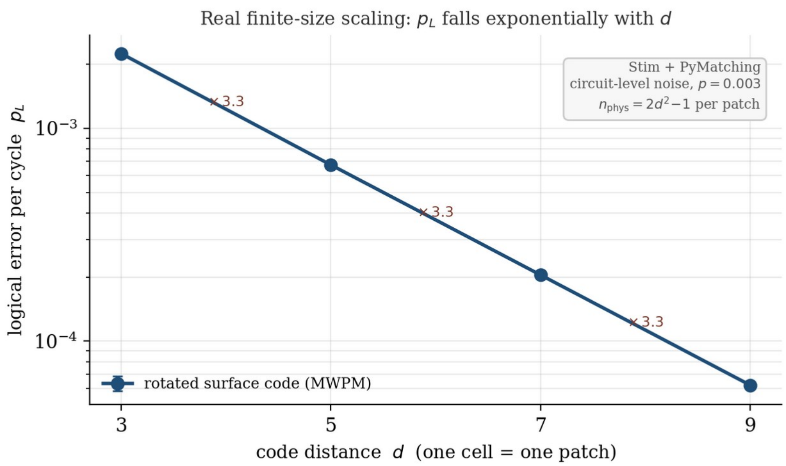

Figure 2. *Real finite-size scaling (Stim + PyMatching). Per-cell logical error falls ~3.3× per +2 distance — but for fixed distance it is a constant, independent of how many cells the machine has.*

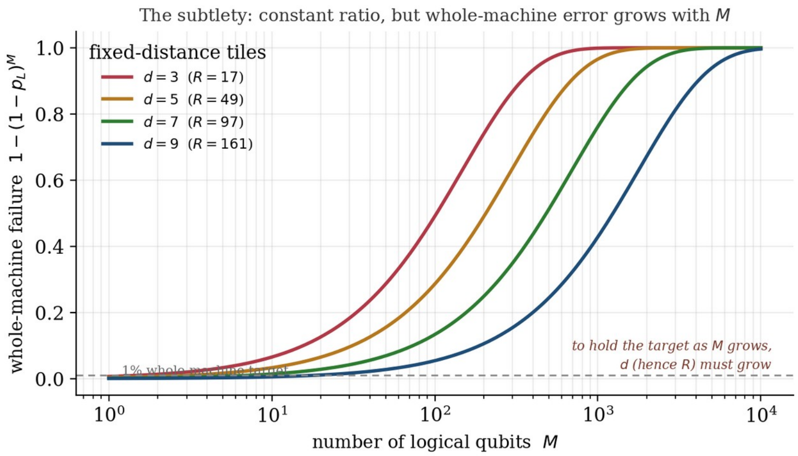

Figure 3. *From the same measured numbers: each fixed-distance tile has a constant ratio R, yet whole-machine failure climbs with the number of logical qubits. Holding a fixed target forces the distance — hence the ratio — to grow.*

So the decisive test is a finite-size stability fit, with an explicit pass criterion — and an honest classification when the error rises but stays under target:

$$\text{fit } p_L(N_{\log}) = \beta_0 + \beta_1 N_{\log} : \quad |\beta_1| \le \delta_\beta \Rightarrow \text{ stable;} \quad \text{rising but} \le \text{target} \Rightarrow \texttt{CONSTANT\_RATIO\_WITH\_BOUNDED\_ERROR\_ONLY}, \text{ not } \texttt{FULL\_NOISE\_INDEPI}$$

The resolution that makes overhead genuinely constant at fixed error is a code family with constant rate and a distance that grows with the code size (good qLDPC) — which cannot come from independent fixed tiles, and which the native admissible search space must deliver. The surface-code run here demonstrates the methodology and the subtlety; it is not a simulation of the corpus's specific architecture.

6. Noise channels, estimator, and replay protocol

The declared channels are explicit — depolarizing and biased Pauli, with measurement, idle, and reset terms in the circuit-level model:

$$\mathcal{E}_{\text{dep}}(\rho) = (1-p)\rho + \tfrac{p}{3}\big(X\rho X + Y\rho Y + Z\rho Z\big), \qquad \mathcal{E}_{\text{bias}}(\rho) = (1 - p_X - p_Y - p_Z)\rho + \sum_{\mu \in \{X,Y,Z\}} p_\mu\, \sigma_\mu \rho\, \sigma_\mu$$

The logical-error estimator is a Monte Carlo failure fraction with its binomial standard error; rare-event points use a binomial confidence interval, and every point reports shots, failures, seed, and decoder:

$$\hat{p}_L = \frac{F}{N_{\text{shots}}}, \qquad \text{SE}(\hat{p}_L) = \sqrt{\frac{\hat{p}_L\,(1-\hat{p}_L)}{N_{\text{shots}}}} \qquad (F = \text{logical failures})$$

and the decisive gate is that this estimator does not grow with system size in the claimed regime:

$$\frac{d}{dN_{\log}}\, p_L(N_{\log}) \approx 0 \qquad \text{(noise stability ˇ2014 Gate Q3, the decisive test)}$$

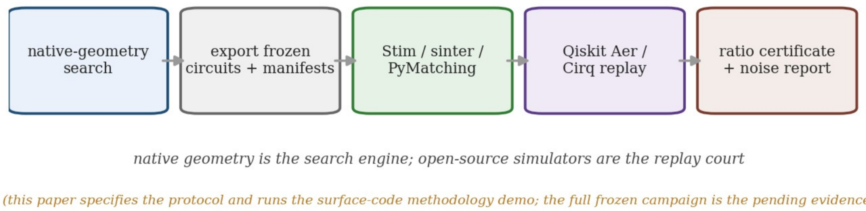

Figure 4. *The replay protocol, not self-graded. Native geometry proposes; Stim / sinter / PyMatching and an independent second simulator dispose. Exported circuits, detector error models, decoder configs, noise manifests, seeds, command lines, version pins, raw logs, and a reproduction script are all required.*

7. Blocked-value ledger, and verdict

Exactly what is in hand, what is pending, and what is blocked — so no reader mistakes the proven counting theorem for a validated overhead number:

Quantity	Status	Note
Surface-code finite-size data (d = 3,5,7,9)	`AVAILABLE_REAL`	Stim + PyMatching; included
Numeric ratio R_star = n(z)/k(z)	`BLOCKED_BY_MISSING_FROZEN_VALUE`	no ratio certificate; symbolic
Architecture-specific p_L(N_log) curve	`PENDING_REPLAY`	demo is methodology only
Search-space collapse factor	`BLOCKED_BY_MISSING_SEARCH_MANIFEST`	no count manifest frozen
Correlated / temporally-correlated noise	`CORRELATED_NOISE_NOT_SIMULATED`	sensitivity run owed
Second-simulator (Qiskit / Cirq) replay	`PENDING_REPLAY`	REPLAY_LIMITATION where unrepresentable
Universal optimality (F_13D = F_std)	`UNIVERSAL_OPTIMALITY_NOT_CLAIMED`	disciplined subspace only

> **Verdict:** `QC_FULL_MATH_COMPLETE_PENDING_REPLAY`. The theorem is stated and proven both ways with every resource counted, the boundary terms and finite-size correction are explicit, the noise channels and estimator are written, and the equivalence theorem holds with `UNIVERSAL_OPTIMALITY_NOT_CLAIMED`. A real surface-code simulation establishes the methodology and shows why noise stability — not counting — is the decisive remaining question. What is pending is the frozen, replayed campaign on the admissible architecture and the ratio

> certificate; no specific physical-to-logical ratio is claimed, and the constant remains R_star.

Notes and sources

[1] The surface-code finite-size data are real: rotated-surface-code memory circuits at $d = 3, 5, 7, 9$ under circuit-level depolarizing noise at $p = 0.003$, generated with Stim and decoded with PyMatching (minimum-weight matching), 0.4–2.0 million shots per point; per-cycle logical error 2.25×10^{-3}, 6.74×10^{-4}, 2.05×10^{-4}, 6.21×10^{-5} for $d = 3, 5, 7, 9$ (suppression $\Lambda \approx 3.3$ per +2 distance). Simulation script and results are included. The run demonstrates the methodology and the union-bound subtlety; it is not a simulation of the corpus's specific architecture.

[2] The constant-rate counting theorem is elementary; the distinction between constant rate and good codes (constant rate with d ~ n^α, giving O(N) physical qubits to protect N logical at fixed error) is the subject of the quantum-LDPC literature (e.g. Panteleev–Kalachev and subsequent constructions).

[3] The corpus's own quantum-error-correction architecture (the $K_6 = SU(3)/T^2$ coset as the admissibility chamber of a designed-and-simulated design) is carried in Paper II ('Forces') §16 at SIMULATED grade, with a reported physical-to-logical memory overhead floor presented as an honest engineering floor competitive with published qLDPC overheads, walled off from the physics. No specific ratio is claimed here: physics.magflowmeters.com

[4] Companion: the base constant-overhead note and the rest of the consistency series (the four-force and quantum-gravity notes, the synthesis, the particle-ontology note, and the metric-interface toy).