An autonomous four-dimensional quantum memory from a time-independent Toom Lindbladian

Seth Douglas Nidhal Mghirbi

November 22, 2026

Abstract

We prove exponential storage for the time-independent four-dimensional Toom dissipative toric code under any sufficiently weak additive finite-range Markovian Lindblad noise generator, including coherent Hamiltonian terms and non-unital dissipation such as amplitude damping. From an encoded start, a decoder reading only the current syndrome has worst-case fidelity at least $1 - Ce^{-cL}$ at every $0 \le t \le e^{cL}$, in both CSS sectors and for all six logical qubits entangled with an untouched reference. The coherent threshold applies to the amplitude $|h|$, not h^2. For such additive noise this realizes the fully autonomous memory setting of Drechsler, Oreg and Aharonov, with translation invariance when the noise is homogeneous. The correction generator is exact; perturbations of its own jump operators and general non-Markovian baths remain open. Separate stochastic-Pauli theorems give explicit syndrome structure and allow spacetime local stochastic bounded-range multi-qubit and adaptive faults on unit time cells (one inverse correction rate), with a noise law independent of the logical input. On $\mathbb{Z}^4$, the independent-noise syndrome process has exponential component-diameter tails uniformly in time from empty syndrome. The corner convention is related to Pastawski, Clemente and Cirac's rule by $x \mapsto -x$ and a uniform X/Z relabelling. The general proof combines Poisson clearance and Toom contours with a coherent protected-region estimate, exact mixed differences and dense-set counting. The thresholds are non-explicit and astronomically small, and the size requirements are astronomically large.

Claude Opus 5.5, GPT-6.1 Sol, GPT-6.2 Sol and GPT-6.2 Astra assisted with this work; the authors take responsibility for every claim.

Contents

1 Introduction **3**

1.1 Relation to existing work 4

1.2 What the probability estimate controls 5

1.3 Why overlapping comparisons are needed 5

1.4 Organization and conventions 5

2 Model and statements **6**

2.1 Cubical code and its two sectors 6

2.2 The upper-corner rule 6

2.3 Time-independent generator and noise . . . 7
2.4 Readout and fidelity convention . . . 7

3 Local estimates without synchronized updates 8
3.1 Confinement and a clearance schedule . . . 9
3.2 Spatial blocks and the local test . . . 10
3.3 Five predecessor regions . . . 12

4 A monotone comparison and localized contours 13
4.1 The auxiliary field and the slab proxy . . . 13
4.2 The deterministic contour input . . . 14
4.3 Five charges, edge count and dependent sinks . . . 14
4.4 Localized spacetime certificates . . . 15
4.5 From certificates to a centered path estimate . . . 17

5 Overlapping comparison windows 18

6 A fixed current-syndrome decoder 20
6.1 Filling a small component . . . 20
6.2 Invariance across a physical face flip . . . 21

7 Quantum evolution and worst-case fidelity 22

8 The process on the infinite lattice 23
8.1 Graphical existence . . . 23
8.2 Uniform component tails . . . 24

9 Spacetime local stochastic Pauli noise 25
9.1 Lindbladian correction, trajectory-defined noise . . . 25
9.2 Noise models included in the hypothesis . . . 27
9.3 Geometry and a joint coarse-bin cost . . . 27
9.4 Joint costs for localized witness families . . . 29
9.5 Unconditional overlapping windows . . . 31
9.6 Pauli branches and current-syndrome recovery . . . 32

10 Additive finite-range Markovian Lindblad noise 33
10.1 Noise strength and theorem . . . 33
10.2 A local Pauli normal form and completely positive interpolation . . . 34
10.3 Intrinsic slab violations and protected regions . . . 36
10.4 Exact linear histories and a one-leg filter . . . 37
10.5 Mixed differences and joint diamond cost . . . 38
10.6 Walk extraction on the finite torus . . . 39
10.7 Dense subsets and the superset polynomial . . . 41

10.8 Linear readout and passage to the exact semigroup 43

11 Scope, parameter choices and limitations **44**

11.1 Order of constants 44

11.2 A wrapping syndrome that does not clear 44

11.3 What is and is not implied 45

1 Introduction

A local correction mechanism can be time independent even when the events by which it acts are random. The question is whether uncoordinated local attempts can suppress the syndrome structures that eventually produce a logical error. In four dimensions the toric code offers a particularly concrete test: syndrome edges form even-degree loops, and a conditional plaquette flip can move their upper corners towards smaller coordinate height.

Pastawski, Clemente and Cirac proposed a dissipative realization of this four-dimensional Toom mechanism and supported its memory properties numerically [13]. Drechsler, Oreg and Aharonov proved synchronous and Floquet memory theorems and asked whether a time-independent Lindbladian, including that proposal, has a nonzero memory threshold [8, Sections 1.5.3 and 1.8]. Their own physical noise model is independent single-qubit Pauli noise: Definition 2.25 gives the sum of one-qubit jump generators, Definition 2.37 the product Pauli layer, and Definition 7.7 the corresponding Floquet noise assumption [8]. Theorem 2.1 therefore answers their time-independent question in their model.

Theorem 9.2 goes beyond that model: the exponential lifetime and fidelity bound also hold under spacetime local stochastic Pauli noise on unit time cells, allowing bounded-range multi-qubit and adaptive faults. The noise law must be independent of the hidden logical input and its reference. For this extension the physical description is *Lindbladian correction, trajectory-defined noise*. Our strongest result, Theorem 10.1, proves exponential storage under any sufficiently weak *additive* finite-range Markovian Lindblad noise generator, including coherent Hamiltonian terms and non-unital noise such as amplitude damping. It controls all six logical qubits and an untouched reference. The coherent threshold is on the amplitude $|h|$, not h^2. Drechsler, Oreg and Aharonov reserve *fully autonomous* memory for protection against general local noise by a time-independent, local, translationally invariant Lindbladian [8, Section 1.5.1, Table 1]. Theorem 10.1 meets that definition for additive finite-range Markovian noise, with translation invariance for homogeneous noise; its bound also permits spatially inhomogeneous noise.

In Section 1.4.1 they describe the setting outside thermal equilibrium as having "no known rigorous proof of a fully autonomous quantum memory" [8]. The thermal precedent is the four-dimensional toric code under Davies dynamics studied by Alicki, M. Horodecki, P. Horodecki and R. Horodecki [2]. Here the added noise generator is required to generate CPTP evolution and the correction remains exact. Perturbations of the correction's own jump operators, such as an imperfect syndrome projector, remain outside the result, as does a general non-Markovian bath. The former belongs to the broader superoperator-stability question in Drechsler–Oreg–Aharonov Section 1.8 [8].

The independent-noise conclusion has two related forms. On the finite four-torus a single, continuously evolving device retains its encoded information for an exponential time, with an exponentially small worst-case fidelity error. On the infinite lattice, the syndrome process started

empty has exponential component-diameter tails at any specified location, uniformly over all physical times. The infinite statement concerns syndrome geometry; it does not assign a logical qubit to the infinite lattice.

1.1 Relation to existing work

The directed geometry and the contour method originate with Toom [16]. Bennett and Grinstein studied the stabilizing role of irreversibility in local stochastic dynamics [4]. Bramson and Gray developed renormalization comparisons with growth and erosion processes [5]; Gray proved continuous-time stability results for attractive spin systems [10]. These works establish a substantial classical background for asynchronous stability.

The present syndrome process is not an attractive spin system under edgewise inclusion. An enabled plaquette operation removes the two tested edges and toggles the opposite two; an opposite edge can disappear if it was already marked. Thus inclusion of edge sets does not give the order needed for a direct attractive-process comparison. We instead construct an auxiliary monotone process on spatial blocks and verify its implications for every admissible starting syndrome in a slab.

Chakraborty and Lee use the same kind of slab comparison to study finite-range monotone continuous-time dynamics [6, Theorem 1 and Appendix B]. The local comparison here is established directly from the parity of syndrome vertices and the waiting times of the physical Poisson clocks. Swart, Szabó and Toninelli provide the deterministic contour construction and its enumeration [15, Theorem 22 and Lemmas 38–40]. We apply that construction to the block process, with a separate calculation for the finite-range dependence of the bad-bin field.

For quantum codes, Dennis, Kitaev, Landahl and Preskill developed the homological account of topological quantum memory and discussed local correction in four dimensions [7]. From Ahn's thesis we use only the five linear forms in Eq. (7.17) [1]; their required inequalities are checked directly in Section 4. No dynamical or threshold estimate from that thesis is an input to this proof. Kubica and Preskill proved threshold results for synchronous cellular-automaton sweep decoders [11]. Ponarovsky treats local stochastic Pauli noise with synchronous sweeps and noiseless bounded-depth classical processing [14, Sections 1.1–1.3, Theorems 6.3 and 7.3]. That paper also notes gaps in Ahn's four-dimensional Toom proof and points to a companion discussion [14, Section 1.1]. Its timing and processing resources differ from independent Poisson correction attempts.

Kumar and Weimer study a two-dimensional dissipative toric code with a dynamical classical field [12]. Dünnweber, Styliaris and Trivedi give a two-dimensional autonomous cellular-automaton construction and outline a continuous-time realization [9]. These are related constructions, not inputs to the estimates below. Joint charging of fault tuples is standard in local-stochastic fault tolerance; the correlated extension combines it with the asynchronous clearance, contour and overlap arguments.

We credit the quantum dissipative proposal and its numerical evidence to Pastawski, Clemente and Cirac [13], and the synchronous and Floquet theorems and the explicit time-independent question to Drechsler, Oreg and Aharonov [8]. The method used here is standard in kind: a block comparison, Toom contours, a Peierls estimate and a homological decoding argument. Its application to the non-monotone four-dimensional toric-code syndrome dynamics appears new.

1.2 What the probability estimate controls

We prove, in either CSS sector, the history bound

$$\begin{aligned}&\mathbb{P}\{\text{a syndrome component of diameter } \geq L/100 \text{ occurs by time } t\}\\&\qquad \leq C_0(1+t)e^{-\beta L} \qquad (t \geq 0).\end{aligned} \tag{1.1}$$

The constants and the lower bound on L are uniform over all allowed spatial choices of the noise rates. The inequality is understood with its right side replaced by its minimum with one. It is a statement about all actual jumps up to a deterministic horizon, rather than only the syndrome at regularly spaced observations.

A fixed filling of each small syndrome component defines a decoded homology class. A single plaquette flip cannot change this class when the components immediately before and after the flip are small. Equation (1.1) therefore controls the entire decoded history. Applying this argument in both sectors gives an event on which every code state, including a state entangled with a reference, is recovered exactly. The worst-case fidelity bound follows from the probability of that event.

1.3 Why overlapping comparisons are needed

A short backward dependency cone on a large torus lifts injectively, date by date, to the integer lattice. A contour in such a cone has distinct physical sinks, so each sink can be charged to a rare collection of Poisson increments. Over a much longer cone, distinct lifted sinks might be copies of the same physical bin. Multiplying their costs would then be invalid.

We avoid this repetition by estimating two comparisons initialized one epoch apart. They use the same maps during their overlap. If they still disagree at its end, a path of zeros of the older comparison crosses an entire epoch in time. Localized spacetime witnesses make such a path exponentially unlikely while the older comparison is still young enough to lift. Agreement can then be propagated from one epoch to the next. The comparisons are functions of the marks; they do not restart the physical memory.

There is a genuine reason to retain the clean-initialization hypothesis. Some wrapping, even-degree syndromes are frozen under every noiseless correction attempt. Consequently the argument does not assert uniform contraction from arbitrary torus syndromes. Section 11 gives an explicit example.

1.4 Organization and conventions

Section 2 specifies the generator, fidelity convention and statements. Section 3 proves the asynchronous clearance and the uniform local mark test. Section 4 constructs localized certificates and proves the centered path estimate. Section 5 establishes (1.1) by overlapping windows. Sections 6 and 7 construct the readout and complete the quantum theorem. Section 8 gives the infinite-volume construction and its uniform estimates. Section 9 states and proves the correlated-noise extension, with its time-grid dependence and joint witness costs. Section 10 proves the strongest, additive Markovian-noise theorem by retaining coherent cross terms, resumming optional faults, and counting dense unions of witnesses.

Chains and boundaries have coefficients in $\mathbb{F}_2$. A syndrome component means a component of marked edges joined at common endpoints. Diameters are ambient sup-metric diameters of their

vertex sets, using periodic coordinate distances on a torus. Constants denoted by C may be enlarged from one inequality to the next. Every fixed block parameter is chosen independently of L.

2 Model and statements

2.1 Cubical code and its two sectors

Let K_L be the cubical cellulation of $(\mathbb{R}/L\mathbb{Z})^4$, with vertices $\mathbb{Z}_L^4$ and unit coordinate cells. Write $C_j(K_L;\mathbb{F}_2)$ for its cellular j-chains and ∂_j for the boundary. There is a physical qubit on each plaquette, so the number of qubits is $6L^4$. For each edge e and cube c, set

$$A_e = \prod_{p \supset e} X_p, \qquad B_c = \prod_{p \subset \partial c} Z_p. \tag{2.1}$$

These checks commute: an edge and a cube have either zero or two common incident boundary plaquettes. Their joint $+1$ eigenspace is the code space. A Z-error is a chain $E \in C_2(K_L;\mathbb{F}_2)$, with syndrome $s = \partial_2 E$ on the edge checks. Since $\partial_1\partial_2 = 0$, every syndrome vertex has even degree. A closed residual Z-chain is trivial on the code precisely when it belongs to $\operatorname{im}\partial_3$.

The dual cubical cellulation K_L^* identifies each primal plaquette p with a dual plaquette p^*, and each primal cube with a dual edge. An X-error is consequently a dual 2-chain with a dual edge syndrome. The two descriptions have the same geometry. The dimension of $H_2(K_L;\mathbb{F}_2)$ is $\binom{4}{2} = 6$; equivalently the coordinate 2-tori give six independent membrane classes. One may see this from the product cellular complex of four circles: its mod-two differentials vanish and its degree-two cells are indexed by pairs of circle factors. Thus the code carries six logical qubits.

2.2 The upper-corner rule

On an integer lift, denote by e_i the coordinate unit vectors and write

$$\mathrm{ht}(v) = v_1 + v_2 + v_3 + v_4.$$

For $i < j$, the face with upper corner v has vertices

$$v, \quad v - e_i, \quad v - e_j, \quad v - e_i - e_j.$$

Its tested edges are $[v - e_i, v]$ and $[v - e_j, v]$. A correction attempt flips that plaquette exactly when both tested edges belong to the syndrome. All six face types at every upper vertex have independent rate-one Poisson attempt clocks. Ineffective rings remain part of the graphical construction. On the dual lattice use the same upper-corner convention in dual coordinates.

This is the convention in Drechsler–Oreg–Aharonov, Definition 6.2 and Eq. (169) [8]. In the notation of Pastawski–Clemente–Cirac, Eq. (33), the tested sides leave the lower corner [13]. The reflection

$$\mathfrak{r} : \mathbb{Z}_L^4 \longrightarrow \mathbb{Z}_L^4, \qquad \mathfrak{r}(x) = -x \tag{2.2}$$

takes our face with upper vertex v to their face with lower vertex $-v$, and takes its tested incoming edges to their outgoing lower sides. The same applies to the dual lattice after choosing its coordinate origin. Check Pauli names are conventional: their edge checks are Z-type, whereas ours are X-type, so identifying the two displayed stabilizer conventions also exchanges X and Z uniformly,

implementable by a Hadamard on every qubit. Reflection preserves incidence, locality and the independent-clock laws. This identifies the geometric rule up to reflection and the stated Pauli convention.

2.3 Time-independent generator and noise

For a primal plaquette p with tested edges e, e', and its dual plaquette with tested dual edges corresponding to cubes c, c', put

$$\Pi_p^X = \frac{1-A_e}{2}\frac{1-A_{e'}}{2}, \qquad \Pi_p^Z = \frac{1-B_c}{2}\frac{1-B_{c'}}{2}. \tag{2.3}$$

Products are unambiguous because all checks commute. Define $\mathcal{D}[J](\rho) = J\rho J^\dagger - \frac{1}{2}(J^\dagger J\rho + \rho J^\dagger J)$. The model is

$$\mathcal{L}_L = \sum_p (\mathcal{D}[Z_p\Pi_p^X] + \mathcal{D}[X_p\Pi_p^Z]) + \sum_p \sum_{P\in\{X,Y,Z\}} \gamma_{p,P}\mathcal{D}[P_p], \qquad 0 \le \gamma_{p,P} \le \eta. \tag{2.4}$$

The coefficients are constant in time. Homogeneous $\gamma_{p,P}$ give translation invariance. The support of each jump has a bounded diameter independent of L. The Hamiltonian is zero.

The noise clocks for all pairs (p, P), and the two families of correction attempt clocks, are mutually independent. A noise ring applies its single-qubit Pauli without a syndrome test. In the primal Z-error marginal the noise rate is $\gamma_{p,Z} + \gamma_{p,Y} \le 2\eta$, and in the dual X-error marginal it is $\gamma_{p,X} + \gamma_{p,Y} \le 2\eta$. A Y mark belongs to both marginals; their independence is neither asserted nor needed.

Correction rate one fixes the unit of physical time. If all correction attempt rates instead equal $\lambda > 0$, rescaling time by λ replaces $\gamma_{p,P}$ by $\gamma_{p,P}/\lambda$. All statements below then apply with $\eta = \max_{p,P} \gamma_{p,P}/\lambda$ and horizon $\lambda^{-1}e^{cL}$.

2.4 Readout and fidelity convention

Choose a code encoding isometry V_L for the six logical qubits. For the designated qubit Q, prepare the other five logical qubits in any fixed density matrix σ. Write

$$\mathcal{E}_{\mathrm{enc}}(\rho) = V_L(\rho \otimes \sigma)V_L^\dagger.$$

The recovery $\mathcal{R}_L$ measures the current commuting checks, applies fixed compatible primal and dual syndrome fillings, inverts the encoding, and discards the other five logical qubits. Section 6 defines the fillings on every attainable syndrome. Define

$$\mathcal{N}_{L,t} = \mathcal{R}_L \circ e^{t\mathcal{L}_L} \circ \mathcal{E}_{\mathrm{enc}}, \qquad F_{\min}(L,t) = \min_{\|\psi\|=1} \langle\psi|\mathcal{N}_{L,t}(|\psi\rangle\langle\psi|)|\psi\rangle. \tag{2.5}$$

This is the squared-overlap convention of Drechsler–Oreg–Aharonov, Definition 2.19. Their Definition 2.20 calls the system an ϵ-memory on an interval I when $\inf_{t\in I} F_{\min}(L,t) \ge 1-\epsilon$ [8]. The recovery in this definition is fixed as a function of the measured syndrome. It needs no earlier syndrome record.

Theorem 2.1 (Finite-volume quantum memory)**.** *There exist constants $\eta_0, c, C > 0$, an integer L_0, and for each integer $L \geq L_0$ a syndrome-only recovery $\mathcal{R}_L$, such that the following holds. For every set of constant rates $0 \leq \gamma_{p,P} \leq \eta_0$, every choice of designated logical qubit and fixed state σ of the other five qubits, and every $0 \leq t \leq e^{cL}$, the dynamics* (2.4)*, initialized by $\mathcal{E}_{\mathrm{enc}}$, satisfies*

$$\inf_{0 \leq t \leq e^{cL}} F_{\min}(L, t) \geq 1 - Ce^{-cL}. \tag{2.6}$$

The constants are uniform over these rates and inputs. In fact, for every pure state $|\Psi\rangle_{QR}$ with an arbitrary reference R,

$$\langle \Psi | (\mathcal{N}_{L,t} \otimes \mathrm{id}_R)(|\Psi\rangle\langle\Psi|) |\Psi\rangle \geq 1 - Ce^{-cL}$$

through the same interval. The recovery can be chosen independently of the rates, t, σ, and the input state.

The proof is completed in Section 7. A single event will recover all six logical qubits, which explains the stated uniformity in the unused logical state. The restriction to an encoded start includes the reference-entangled version.

Theorem 2.2 (Infinite-volume syndrome stability)**.** *For either sector on $\mathbb{Z}^4$, the independent-clock graphical process started with empty syndrome is well defined. There exist $\eta_0, C, \mu > 0$ such that, for every choice of constant single-qubit Pauli rates bounded by η_0, every vertex x, every deterministic $t \geq 0$, and every real $d \geq 0$,*

$$\mathbb{P}\{x \text{ belongs at time } t \text{ to a syndrome component of diameter } \geq d\} \leq Ce^{-\mu d}. \tag{2.7}$$

The same bound, after a fixed enlargement of C, holds if the component and its membership of x are observed at any common instant of the slab of the proof containing t. Moreover, writing $\mathbb{P}_\gamma$ for an allowed family of rates,

$$\lim_{\eta \downarrow 0} \sup_{\substack{0 \leq \gamma_{p,P} \leq \eta \\ x \in \mathbb{Z}^4,\ t \geq 0}} \mathbb{P}_\gamma\{x \text{ is occupied at time } t\} = 0.$$

Diameter is the ambient lattice sup diameter. The assertions are marginal statements for a named sector.

Section 8 proves this theorem, including existence of the process. At $t = 0$ the syndrome is empty. In statements about a whole slab we use the slab $[0, T]$ for $t = 0$, and otherwise $((k-1)T, kT]$ when $(k-1)T < t \leq kT$.

Theorems 2.1 and 2.2 have the independent single-qubit Pauli scope stated above. Theorem 9.2 treats correlated stochastic Pauli faults. The strongest statement, Theorem 10.1, covers arbitrary sufficiently weak additive finite-range Markovian Lindblad noise, including coherent and non-Pauli terms, with the exact correction. Perturbations of the correction's own jump operators and general non-Markovian baths remain outside that theorem. The estimates yield existence of positive constants, rather than a numerically useful threshold or a practical minimum size.

3 Local estimates without synchronized updates

We work in one sector. Every correction ring toggles the boundary of one plaquette, conditionally on its two incoming edges, and a sector noise mark toggles the same boundary without a test. Write

a vertex as occupied when it is incident to at least one marked edge. The following deterministic facts refer to the actual sequence of rings.

3.1 Confinement and a clearance schedule

Lemma 3.1 (Geometric confinement)**.** *For a finite marked edge set on $\mathbb{Z}^4$ evolving with no noise under any sequence of upper-corner correction attempts, the coordinate box of occupied vertices cannot enlarge and their maximum height cannot increase. Each vertex degree modulo two is preserved. The confinement assertions hold even if the starting degrees are odd.*

Proof. An enabled face at v has its tested edges present. Consequently $v, v - e_i, v - e_j$ are already occupied. The remaining corner $v - e_i - e_j$ lies in their coordinate box and has smaller height than v. All vertices on potentially new edges are therefore in the existing box and below or at its maximum height. Toggling a plaquette boundary changes the degree of each of its four vertices by an even integer modulo two. Induction over the attempts proves every assertion. Ineffective attempts change nothing. □

Fix an integer H and a vertex x with $\mathrm{ht}(x) \le H$. The northeast scheduling cone is

$$\mathcal{C}(x, H) = \{x + a : a \in \mathbb{Z}^4_{\ge 0},\ \mathrm{ht}(x + a) \le H\}.$$

At each vertex v, let $F_v(s)$ be the time of completion of the following subsequence of its actual clocks after time s: wait for the six pairs $i < j$ in lexicographic order, and repeat that list. Other rings, including rings of these clocks while waiting for a different pair, are allowed. For a fixed realization each next-ring function, and hence F_v, is nondecreasing.

Set $S(v) = 0$ above height H, and recursively set

$$S(v) = F_v\left(\max_{1 \le i \le 4} S(v + e_i)\right) \qquad (v \in \mathcal{C}(x, H)). \tag{3.1}$$

Times here are measured relative to the beginning of a clean interval.

Lemma 3.2 (Poisson clearance)**.** *Suppose a clean process has initial maximum occupied height at most H, even degree at every vertex of $\mathcal{C}(x, H)$, and all faces incident to that cone available for correction. This includes a finite cut if the cone and its incident faces lie in its interior. After time $S(x)$, the vertex x has no incident marked edges, permanently. With independent rate-one clocks and $m = H - \mathrm{ht}(x) + 1$,*

$$\mathbb{P}\{S(x) > t\} \le 4^m 2^{12m} e^{-t/2} \qquad (t \ge 0). \tag{3.2}$$

The event $S(x) \le t$ reads only the clocks up to t.

Proof. An incoming edge at a vertex above height H is always absent by Lemma 3.1. Suppose the four vertices $v + e_i$ have permanently lost their incoming edges by time s. The four outgoing edges from v are then absent. An incoming edge at v can be affected only by a face with upper corner v, or one of those four upper neighbours. A face based at an upper neighbour cannot subsequently be enabled, because its incoming edges are absent.

The number of incoming marked edges at v is now even by parity. Every enabled face based at v removes two such edges; none can create an incoming edge there. In one full pass through

the six pairs, any fixed remaining pair either rings while still present, or some removal has already happened. Thus a pass removes at least two whenever at least two remain. There are at most four incoming edges. Two passes remove them all. Thereafter no tested pair can be enabled at v, and the upper neighbours cannot recreate an incoming edge. Descending induction on height proves the deterministic claim using (3.1). Boundary parity in a cut is unnecessary away from the indicated cone.

For the probability bound, monotonicity of F_v implies $F_v(\max_i s_i) = \max_i F_v(s_i)$. Expanding (3.1) therefore expresses $S(x)$ as the maximum of serial completions along northeast paths of m vertices, read from the largest height downwards. There are at most 4^m such paths. Along any specified path the upper vertices are distinct. Its serial completion time is a sum of $12m$ independent rate-one exponentials. Indeed the start time for the clocks at each successive vertex depends only on clocks at distinct earlier vertices, and the next waiting times at that vertex have the exponential law by independent increments and the strong Markov property. For a rate-one exponential W, $\mathbb{E}e^{W/2} = 2$. The exponential Markov inequality followed by a path union bound gives (3.2).

Finally the recursive test for completion by t either finds all required subsequence rings by t, or fails to do so. It never needs to inspect a ring after t. This proves its measurability. □

Corollary 3.3 (An isolated finite syndrome)**.** *An even syndrome initially confined to an integer box of maximum coordinate side length d, with no noise, becomes empty by time t except with probability at most*

$$\min\{1, (d+1)^4 (4 \cdot 2^{12})^{4d+1} e^{-t/2}\}.$$

Every correction stays in that box. If the initial error chain is supported in the box, its final closed residual is a boundary. The same assertions hold for a torus box whose one-cell collar lifts injectively, with independent torus clocks restricted to that box.

Proof. Choose H to be the initial maximum height. For any vertex in the box, $H - \mathrm{ht}(x) \le 4d$. Apply Lemma 3.2 and sum over its at most $(d+1)^4$ vertices. Confinement gives the support assertion. A closed 2-chain in a cubical box bounds there; an explicit chain contraction is provided in Section 6.

For the torus assertion transfer the clocks in the box and its collar to its integer lift, and add independent clocks outside that collar. Confinement makes every enabled face stay inside the original box, so projection of this infinite-lattice construction agrees with the confined torus trajectory. The auxiliary scheduling cone may extend beyond one period. Its additional clocks are independent extension clocks, rather than repeated torus clocks. The estimate thus uses distinct independent vertices along each path and remains valid. □

The last distinction is useful for isolated components but will not justify arbitrarily deep lifts of interacting torus histories. The short-window construction below is what guarantees distinct physical defects in a contour.

3.2 Spatial blocks and the local test

Fix an integer $B \ge 10$. For L sufficiently large define

$$M = \lfloor L/B \rfloor, \quad b = L/M, \quad T = 1000B, \quad m_f = 1000T, \quad R = 2000100. \tag{3.3}$$

The coarse lattice is $\mathbb{Z}^4_M$. In a consistent lift its blocks are

$$Q_z = \{v \in \mathbb{Z}^4 : bz_i \le v_i < b(z_i + 1), \quad 1 \le i \le 4\}.$$

Here $B \le b < 2B$. Translation by M in a block coordinate translates the physical coordinate by exactly L. Throughout this section assume $M > 10(R + 5)$. A bin (z, k) refers to the physical slab $I_k = ((k - 1)T, kT]$. Lift the blocks with $\|d\|_\infty \le R$ relative to z, and let

$$U_z = \bigcup_{\|d\|_\infty \le R} Q_{z+d}.$$

To compare with the device, retain the initial syndrome edges whose two endpoints lie in U_z, delete the other edges, and run the clean dynamics using only faces entirely in U_z. Its clocks are the physical clocks of these faces.

Definition 3.4 (Bad bin)**.** The indicator $J(z, k)$ equals zero when all three conditions below hold, and equals one otherwise:

1. No sector noise mark in I_k occurs on a face with any vertex in U_z.
2. There is no strictly chronological sequence of m_f correction rings on successively vertex-intersecting faces, ending on a face within physical sup distance two of Q_z. Count such sequences in the lifted neighbourhood of Q_z.
3. For $H_z = \lfloor b(z_1 + z_2 + z_3 + z_4 + 4) \rfloor$, every $x \in Q_z$ has $S(x) \le T$, with the schedule of (3.1) started at the beginning of I_k.

Each sequence in condition 2 lies within distance $m_f + 4$ of Q_z, and hence in U_z. Each cone in condition 3 has depth at most $8B + 2$, and its incident faces are interior to U_z. These tests depend solely on marks, not on the starting syndrome.

Lemma 3.5 (Probability and range of dependence)**.** *For every $q > 0$ there exist an integer $B \ge 10$ and $\eta_0 > 0$, independent of L, such that for every allowed set of rates bounded by η_0, all z, k, and all past histories up to $(k - 1)T$,*

$$\mathbb{P}(J(z, k) = 1 \mid \textit{past}) \le q.$$

The test reads only face processes anchored in blocks at coarse distance at most $R + 3$ from z, during I_k. Disjoint sets of such basic increments are independent.

Proof. There are at most $C_R B^4$ faces incident to U_z. The sector noise rate per face is at most 2η, so condition 1 fails with probability at most $2\eta C_R B^4 T$.

A vertex of the four-dimensional cubical lattice is incident to $4\binom{4}{2} = 24$ faces. A face has four vertices, and so meets at most 96 faces at vertices, including itself. There are at most CB^4 choices of a terminal face in condition 2. For each specified list of m_f face labels, the expected number of ordered ring tuples in a slab is $T^{m_f}/m_f!$. This follows by integrating the rate-one product intensity over the time simplex. If a face label repeats, the factorial moment measure of its Poisson process gives the same expression; ordered tuples use distinct times. Consequently condition 2 fails with probability at most

$$CB^4 \frac{(96T)^{m_f}}{m_f!} \le CB^4 (96eT/m_f)^{m_f}.$$

There are at most $(2B+1)^4$ target vertices in condition 3. Lemma 3.2 bounds its failure probability by

$$(2B+1)^4(4\cdot 2^{12})^{8B+2}e^{-T/2}.$$

Combining these estimates gives

$$q_B(\eta) := 2\eta C_R B^4 T + CB^4(96e/1000)^{1000T} + (2B+1)^4(4\cdot 2^{12})^{8B+2}e^{-T/2}. \tag{3.4}$$

The last two terms tend to zero as B tends to infinity: $96e/1000 < 1$, and $8\log(4\cdot 2^{12}) - 500 < 0$. Choose B to make their sum at most $q/2$, and then choose a strictly positive η_0 to make the noise term at most $q/2$.

The face anchors needed for the three tests are in the stated $R+3$ block neighbourhood. All tests use increments within I_k. Those increments are independent of the entire past, even if the initial syndrome was selected from it. The conditional assertion and the independence statement follow. □

3.3 Five predecessor regions

Let

$$\begin{aligned} \Delta &= \{d \in \mathbb{Z}^4 : \|d\|_\infty \le R\}, \\ A_i &= \{d \in \Delta : d_i \le 0\} \quad (1 \le i \le 4), \\ A_5 &= \{d \in \Delta : d_1 + d_2 + d_3 + d_4 \ge 1\}. \end{aligned} \tag{3.5}$$

All five sets are nonempty and $\bigcup_s A_s = \Delta$.

Lemma 3.6 (Uniform predecessor implications)**.** *For every even starting syndrome and every realization with $J(z,k) = 0$, occupation of Q_z at time kT implies, for each $s = 1,\dots,5$, occupation at time $(k-1)T$ of some block Q_{z+d} with $d \in A_s$. Occupation of Q_z at any instant of I_k implies occupation of some block of U_z at its start.*

Proof. We justify agreement of the clean cut and the physical trajectory near Q_z for the entire slab. Trace backwards any differing incident edge. An interior correction ring can create a difference only if a tested edge or the old output edge already differed. Thus a discrepancy arriving from the cut boundary gives a chronological chain of vertex-intersecting faces from that boundary to a face within two of Q_z. Initial interior edges agree, and condition 1 excludes all relevant interior noise. In terms of upper face anchors, consecutive intersecting faces differ by at most one in each coordinate. Endpoint allowances are bounded by a fixed number of cells. The cut boundary is at distance at least $RB - 2$, which exceeds $m_f + 5$. Its final m_f interior rings would contradict condition 2. This argument applies to a discrepancy at any time of the slab. There are almost surely finitely many rings on the finite torus in bounded time.

If no initial block with offset in A_i is occupied, every retained marked endpoint has ith coordinate at least $b(z_i + 1)$. Confinement prevents an endpoint from entering Q_z, for the whole slab. If no initial block in A_5 is occupied, each retained endpoint v satisfies

$$\mathrm{ht}(v) < b(z_1 + z_2 + z_3 + z_4 + 4).$$

Thus its maximum height is at most H_z. All vertices of the scheduling cones retain even initial degree, because they are interior and no incident edge was deleted there. Condition 3 and Lemma 3.2

clear the target vertices by the slab end. Finally, if U_z was initially empty, the clean cut stays empty throughout the slab. Agreement with the physical trajectory proves the two implications by contraposition. □

The monotonicity facts used here concern coordinate extrema, height and the clock-completion functions. They require no ordering of syndromes under inclusion and no synchronized correction rounds. The bins are observations of the original marks; no part of the generator is scheduled or reset at their boundaries.

4 A monotone comparison and localized contours

4.1 The auxiliary field and the slab proxy

For a deterministic slab index $a \geq 0$, initialize $H^a(z, a) = 1$ at every coarse site. For $k > a$ use the same physical bad-bin indicators to define

$$H^a(z,k) = \begin{cases} 0, & J(z,k) = 1, \\ \max\limits_{1 \leq s \leq 5} \min\limits_{d \in A_s} H^a(z+d, k-1), & J(z,k) = 0. \end{cases} \tag{4.1}$$

A one denotes a healthy comparison bit. This map is monotone in its arguments. Also introduce the event, identified with its indicator,

$$P^a(z,k) = \{J(z,k) = 1\} \ \cup \ \{H^a(z+d, k-1) = 0 \text{ for some } d \in \Delta\}. \tag{4.2}$$

Every zero of $H^a(z,k)$ belongs to $P^a(z,k)$. For the physically empty initial syndrome, Lemma 3.6 gives by induction

$$\begin{aligned} Q_z \text{ occupied at } kT &\implies H^0(z,k) = 0, \\ Q_z \text{ visited in } I_k &\implies P^0(z,k). \end{aligned} \tag{4.3}$$

For the induction, an occupied endpoint of a good bin has an occupied predecessor in each A_s. Each of those bits is zero, so every minimum in (4.1) is zero. For the whole-slab assertion, the additional implication in Lemma 3.6 supplies at least one occupied predecessor in Δ. A bad bin belongs to the proxy without any predecessor requirement.

We use the product metric

$$\rho((z,k),(y,\ell)) = \max\{\mathrm{dist}_\infty(z,y), |k-\ell|\} \tag{4.4}$$

on $\mathbb{Z}^4_M \times \mathbb{Z}$, with periodic spatial distance. On $\mathbb{Z}^4 \times \mathbb{Z}$, use the ordinary spatial sup distance. A spacetime path has consecutive distinct sites at ρ-distance at most R. A window has dates $a+1, \ldots, a+D_w$, where on the torus

$$2RD_w < M/2. \tag{4.5}$$

On the infinite lattice any finite depth is permitted.

4.2 The deterministic contour input

We specify the imported combinatorial input so that its use is separate from all probabilistic assumptions. Swart–Szabó–Toninelli [15, Definitions 15–21 and Theorem 22] work with a countable typed dependence graph. A vertex carries a constant-zero, constant-one, or nonconstant monotone map. At a nonconstant vertex there is a finite nonempty set of outgoing predecessors of each charge. The associated map is the maximum, over charges, of the minimum of the predecessor values for that charge. Constants have no outgoing dependence edges. Their Theorem 22 states that a zero of the maximal trajectory has a rooted connected Toom contour.

For clarity, the relevant contour structure consists of sources, sinks and internal vertices carrying one of σ charges. A source emits one edge of every charge and has no incoming edges; a sink receives one of each and has no outgoing edges; an internal vertex has one incoming and one outgoing edge of its charge. The root is a distinguished source. The zero-edge contour consists of a single vertex serving as source and sink. An embedding has distinct sink images, also distinct from all other vertex images, and distinct images of internal vertices of a given charge, counting the root with those internal vertices. These are the injectivity conditions of their Definition 17.

In a present contour the sinks are constant-zero sites and all other vertices are nonconstant sites. Edges leaving the root or an internal vertex follow the allowed predecessors of their own charge. An edge leaving another source may follow a predecessor of any charge. These are the properties of their Definition 21 and Eq. (2.16) that we need. No probability assumption occurs in the contour implication.

In our application a nonconstant vertex at (z, k) has type-s predecessors $(z + d, k - 1)$, $d \in A_s$. All these sets are finite and nonempty, as required by Definition 20 of [15]. A bad bin is constant zero. To use a clean boundary, truncate a root's backward dependency cone at date a, copy the actual maps inside it, and put constant-one maps at date a and everywhere outside, including the future. This finite acyclic recursion has a unique trajectory at the root; in particular it agrees there with the maximal trajectory in Theorem 22. Backward induction on the dates identifies that value with the root of (4.1).

For a torus root, lift this finite cone to $\mathbb{Z}^4 \times \mathbb{Z}$. A section at a fixed date has coordinate widths at most $2RD_w$. Condition (4.5) makes projection injective on that section. Thus distinct lifted sinks at the same date are distinct physical bins. At distinct dates, bins already have disjoint time indices. Constant-one cutoff sites cannot be contour vertices, so every sink is a distinct actual failed test inside the cone. On the infinite lattice the finite cone needs no lift.

4.3 Five charges, edge count and dependent sinks

Use the spatial linear forms of Ahn [1, Eq. (7.17)]:

$$\ell_i(d) = -d_i \quad (1 \le i \le 4), \qquad \ell_5(d) = d_1 + d_2 + d_3 + d_4. \tag{4.6}$$

Their sum is identically zero, and for $d \in A_s$,

$$\ell_s(d) \ge \varepsilon_s, \qquad (\varepsilon_1, \dots, \varepsilon_5) = (0, 0, 0, 0, 1).$$

For any $d \in \Delta$ the respective losses are at most $R, R, R, R, 4R$. All contour edges change time by -1. This checks the step restrictions in [15, Eq. (5.7)], with five charges and positive aggregate advance $\sum_s \varepsilon_s = 1$.

Let a contour have m sinks. The number of sources is also m, by counting the net flux of any one charge. By [15, Lemma 38, Eq. (5.12)], there are equally many edges of every charge; denote their common number by n. This equality also follows directly by summing time differences along the paths of each charge from sources to sinks. Each edge has time difference -1, while the endpoint sum is the same for each charge.

For each charge, $m-1$ edges leave non-root sources. The other $n-m+1$ edges have their prescribed type. Sum $\ell_s(\psi(w)-\psi(v))$ over all charged edges (v,w) of charge s, and over s. Internal contributions cancel; the sum at each source or sink vanishes by $\sum_s \ell_s = 0$. The total is therefore zero. The lower bounds on increments give

$$0 \geq (n-m+1) - 8R(m-1).$$

Thus, with $\alpha = 1+8R$,

$$n \leq \alpha(m-1), \qquad m \geq 1 + n/\alpha. \tag{4.7}$$

This is the specialization of [15, Lemma 40]; the calculation checks its polar-function hypotheses directly.

For the contour enumeration we use the version arXiv:2202.10999v3 of [15]. With $\sigma = 5$ and $\tau = \lceil \sigma/2 \rceil = 3$, Lemma 39, Eq. (5.15), states $N_n \leq n^2 4^{6n} |\Delta|^{5n}$. The sharper odd-charge estimate in the *proof* of that lemma, Eq. (5.28), gives

$$N_n \leq n^2 4^{6n-1} |\Delta|^{5n} \quad (n \geq 1), \qquad N_0 = 1. \tag{4.8}$$

Here N_n counts rooted embeddings up to their equivalence of Definition 19, subject to the step restrictions just verified. Duplicating the last charge in the proof produces six traversals per n, but records the duplicated spatial choices once. The factor 4^{6n-1} is consequently a proof estimate, not the bound stated in the published Lemma 39. Either estimate is exponential in n and would suffice below. Restriction to a deterministic finite cone can only reduce this count.

Dependence of the bad-bin field is handled without importing an independent-defect stability theorem. Put

$$D_{\mathrm{dep}} = (4(R+3)+1)^4. \tag{4.9}$$

For any specified set of m distinct physical bins, join two when they have the same date and spatial distance at most $2(R+3)$. Its maximum degree is at most $D_{\mathrm{dep}} - 1$. Greedy selection gives at least m/D_{dep} bins with disjoint underlying increments. Lemma 3.5 and independence therefore imply

$$\mathbb{P}\{\text{all specified bins are bad}\} \leq q^{m/D_{\mathrm{dep}}}. \tag{4.10}$$

Distinct dates give disjoint Poisson intervals even at the same spatial anchor. For equal dates, the selection ensures disjoint spatial supports. This is where distinct physical sink indices, rather than distinct vertices only on a covering lattice, matter.

4.4 Localized spacetime certificates

Proposition 4.1 (Certificate with adjustable cost)**.** *For every $A > 0$ one can choose $q, B, \eta_0 > 0$ in the order of Lemma 3.5, independent of L, a, D_w, with the following property. For each site $b_0 = (z,k)$ in a window satisfying* (4.5)*, there are events $\mathcal{E}^a(b_0, r)$, $r \in \mathbb{Z}_{\geq 1}$, such that*

$$P^a(z,k) \implies \bigcup_{r \geq 1} \mathcal{E}^a(b_0, r), \qquad \mathbb{P}(\mathcal{E}^a(b_0,r)) \leq e^{-Ar}. \tag{4.11}$$

Each event reads only basic Poisson increments whose pair (anchor block, slab index) lies in the closed ρ-ball $\mathcal{B}_\rho(b_0, r)$. The assertions hold on $\mathbb{Z}^4$ at every finite depth, with the same constants. They also hold conditionally on the entire physical history up to aT.

Proof. If the current bin is bad, use its single failed test. Otherwise a zero predecessor $(y, k-1)$, with $\mathrm{dist}_\infty(y, z) \le R$, is supplied by (4.2). When $k - 1 = a$ there is no such zero. At larger depths apply the finite-cone construction above. Theorem 22 of [15] produces a contour whose m sinks are distinct bad bins and whose edge count obeys (4.7).

Fix $K = 10(R + 3)$. For each $n \ge 0$, use radius $r = K(n + 1)$. Define $\mathcal{E}^a(b_0, r)$ to be the union, over the at most $(2R + 1)^4$ predecessor roots and their candidate contours with n edges per charge in the deterministic finite cones, of the events that the named sinks are bad. At $n = 0$ include the current failed test as well. At other radii use the empty event. The sink events are deliberately larger than the events of actual contour presence; no intermediate good-map conditions need to be read.

A connected contour has $5n$ edges. Along an undirected walk from its root every edge changes space by at most R and time by one, even if that walk reverses the edge direction. Thus each sink lies within $5Rn + R$ spatial units and $5n + 1$ dates of b_0. Testing that sink adds $R + 3$ spatial units and no further date. The radius $K(n + 1)$ contains all those increments. This proves the asserted locality for every candidate in the union.

Equations (4.8)–(4.10) give constants C_0, C_1, depending only on R, such that

$$\mathbb{P}(\mathcal{E}^a(b_0, K(n + 1))) \le C_0 C_1^n q^{(1+n/\alpha)/D_{\mathrm{dep}}} \qquad (n \ge 0). \tag{4.12}$$

The harmless polynomial n^2 has been included in C_1^n. Choose $0 < q < 1$ small enough that both the constant part and the coefficient of n in the logarithm of this bound are at most $-AK$. Then (4.11) follows. The physical mark test realizes this q by choosing B and then η_0.

The union implication follows by selecting the contour furnished by the deterministic theorem, or the direct failed test. All cones and all candidate embeddings are specified without conditioning on their intervening marks. This also proves the measurability claim. Every sink lies after a; all increments used are independent of the history through aT. The same proof therefore gives the conditional bounds. On $\mathbb{Z}^4$ no quotient identifies sink indices, so the argument works at every finite date. □

Witness balls that wrap the torus. The radius $K(n + 1)$ is not restricted by the period M. A spatial ball of radius above $M/2$ wraps the torus and can cover all its spatial sites. This is harmless. Its meaning is the actual closed ball in the periodic product metric, and its support is a set of physical increments, with no repeated copies. Cone injectivity was used before this enlargement to charge the sinks of a single candidate; it is not asserted for the enlarged ball. The volume bound

$$|\mathcal{B}_\rho(b_0, r)| \le (2r + 1)^5$$

still holds, because taking a quotient can only decrease the number of spatial sites relative to that bound. Independence below is used only for balls disjoint as sets on the actual torus times time. Wrapping balls that overlap are never declared independent.

4.5 From certificates to a centered path estimate

Proposition 4.2 (Centered path bound)**.** *Choose A in Proposition 4.1 sufficiently large, and fix the resulting parameters. There exist $C, \kappa > 0$, independent of L, a, D_w, z, k, such that for every anchor b_0 in a permissible window and every real $u \geq 0$,*

$$\mathbb{P}\left\{ \begin{array}{l} \text{a finite path in } P^a \text{ starts at } b_0 \\ \text{and reaches } \rho\text{-distance at least } u \end{array} \right\} \leq Ce^{-\kappa u}. \tag{4.13}$$

All vertices of the path are required to lie in that window. On $\mathbb{Z}^4$ the window can be the entire half-space of dates $k > a$. The bound is also valid conditionally on the physical history through aT.

Proof. Take a finite path realizing the event. Assign to each of its distinct vertices one occurring witness ball from Proposition 4.1. There are finitely many assigned balls. Select disjoint closed balls greedily in decreasing radius. Each discarded ball meets a selected ball of at least its radius. Its center, and indeed the whole discarded ball, lies in the triple enlargement of that selected ball by the triangle inequality. Thus selected triple balls cover the path.

If consecutive path vertices have covering centers c_i, c_j and selected radii r_i, r_j, then

$$\rho(c_i, c_j) \leq 3r_i + 3r_j + R.$$

The graph of these adjacencies connects the covers of the endpoints. Choose a simple chain of selected centers between those covers, possibly a single center. Let its length be j and $s = r_1 + \cdots + r_j$. Endpoint displacements contribute at most $3r_1 + 3r_j$. Summing the displayed adjacency bounds gives

$$u \leq 6s + R(j-1) \leq 7s,$$

because each allowed radius is at least $K > R$. This geometric argument uses the product metric itself and remains valid for wrapping balls.

It remains to count deterministic candidate chains. For fixed ordered positive radii summing to the integer s, the initial center lies within $3r_1$ of the prescribed anchor, so has at most $(6r_1 + 1)^5$ choices. Each subsequent center has at most $(6r_i + 6r_{i+1} + 2R + 1)^5$ choices. These are upper bounds both in the lattice and on the torus times time. Since $r_i \geq K > R$, the product is at most $e^{C_2 s}$ for an absolute C_2. For example, the initial factor costs at most e^{15r_1}; the logarithm of each later factor is at most $20(r_i + r_{i+1})$. Thus one may take $C_2 = 60$.

There are at most 2^s ordered compositions of s. Count only chains whose original closed balls are disjoint. Their witness events read disjoint basic increments, in space and time together, and therefore have joint probability at most e^{-As}. Taking $A > C_2 + \log 2 + 1$ and summing gives

$$\mathbb{P}(\text{path event}) \leq \sum_{s \geq \max\{1, \lceil u/7 \rceil\}} e^{-(A - C_2 - \log 2)s} \leq Ce^{-\kappa u}.$$

The union is over deterministic chains, so no independence is claimed for centers selected after observing the marks. For an unbounded half-space every witnessing path is finite, and the same countable union gives the estimate. Conditional independence of increments after aT proves its conditional form. □

In particular the bound controls temporal paths of auxiliary zeros as well as spatial paths of visited blocks. This simultaneous use of space and time is what permits the window iteration.

5 Overlapping comparison windows

The comparison H^0 dominates the syndrome indefinitely by (4.3), but its full backward cone eventually wraps. We now use only short, clean-start comparison estimates to control that same old field.

Fix all parameters from Proposition 4.2. For sufficiently large M, set

$$h = \left\lfloor \frac{M}{100(R+3)} \right\rfloor, \qquad a_j = jh \quad (j \geq 0). \tag{5.1}$$

An epoch consists of h slabs, and a two-epoch comparison window has depth $2h$. It satisfies $2R(2h) < M/2$, so all results of Section 4 apply within it. We assume $h \geq 1$.

On one common sample of physical marks let $H_j = H^{a_j}$. Each of these auxiliary fields starts all ones at its own date; none restarts the device. Define

$$G_j = \{H_j(z, a_{j+2}) = H_{j+1}(z, a_{j+2}) \text{ for every } z\}, \tag{5.2}$$

$$V_j = \left\{ \begin{array}{l} \text{at no date } k \in (a_j, a_{j+2}] \text{ does } P^{a_j}(\cdot, k) \\ \text{contain a spatially connected set of diameter } \geq M/200 \end{array} \right\}. \tag{5.3}$$

Spatial adjacency here is sup distance one. The diameter is the ambient coarse torus diameter.

Lemma 5.1 (Coalescence over an overlap)**.** *There are constants $C, \kappa > 0$, uniform in j, L and the allowed rates, such that*

$$\mathbb{P}(G_j^c) \leq CM^4 e^{-\kappa h}, \qquad \mathbb{P}(V_j^c) \leq C(2h)M^4 e^{-\kappa M/200}. \tag{5.4}$$

Both estimates hold conditionally on the entire physical history through $a_j T$.

Proof. During the common part of their lifetimes, $H_j \leq H_{j+1}$: the latter is initialized all ones at a_{j+1}, and every later map is monotone or constant. Suppose the two fields disagree at a specified endpoint (x, a_{j+2}). The older bit there is zero and the newer bit is one. They use the same update map. A constant map cannot give different outputs. For a nonconstant map, different outputs imply that some predecessor argument differs, with displacement in Δ. Choose one such predecessor and repeat backwards. This produces a path of disagreements at every integer date back to a_{j+1}, including that date.

Every site on this path is a zero of H_j. Its dates lie strictly after a_j, since $a_{j+1} > a_j$, so every site belongs to P^{a_j}. Consecutive sites have product distance at most R. The endpoint-to-endpoint temporal displacement is h. Proposition 4.2, anchored at the specified final site, bounds this event by $Ce^{-\kappa h}$ in the two-epoch window. Summing over the M^4 possible final spatial sites proves the bound for G_j^c.

If V_j fails, at some one of its $2h$ dates a spatial path in the proxy joins two sites separated by at least $M/200$. The finite coarse torus ensures that a connected set witnessing the diameter contains such a finite path. Apply Proposition 4.2 at its initial spatial site, with $u = M/200$, and sum over M^4 sites and $2h$ dates. This proves the other estimate.

Both comparisons in this argument are clean-start functions of the increments after $a_j T$. Those increments are independent of the history through $a_j T$. Thus precisely the same estimates hold under the stated conditioning. No conditional estimate is used at the middle date of the overlap. □

Lemma 5.2 (Propagation of comparison agreement)**.** *For every integer $n \geq 1$, on*

$$\bigcap_{0 \leq j < n} (G_j \cap V_j)$$

the physical syndrome, started empty, has only components of physical diameter less than $L/100$ at every time in $[0, nhT]$.

Proof. We track which clean-start comparison equals the old comparison. At the beginning of epoch $i \geq 1$ we claim

$$H^0(\cdot, a_i) = H_{i-1}(\cdot, a_i). \tag{5.5}$$

For $i = 1$ this is identity, since $H_{i-1} = H_0$. If it holds at a_i, the two fields use common deterministic maps and remain equal through the following epoch. At a_{i+1}, the event G_{i-1} equates H_{i-1} with H_i. This proves (5.5) inductively for all needed i.

Let W_k be the blocks visited by syndrome vertices at any instant of slab I_k. Equation (4.3) puts W_k in $P^0(\cdot, k)$. In epoch zero this is directly controlled by V_0. In epoch $i \geq 1$, equality (5.5) and its propagation ensure equality of the bits used to form $P^0(\cdot, k)$ and $P^{a_{i-1}}(\cdot, k)$. This includes the predecessor date of the earliest slab in the epoch: it is exactly a_i. Thus W_k is contained in the proxy tested by V_{i-1}. The physical syndrome at that seam can be nonempty; only comparison bits have been equated.

Two adjacent physical vertices belong to identical or adjacent blocks. Further, if their containing blocks have coarse torus distance d, their physical torus distance is at most $b(d+1)$. To verify this inequality, lift their relative block displacement in each coordinate by a shortest representative; their offsets inside the two intervals add less than one block width. This works also for noninteger $b = L/M$.

If a physical component had diameter at least $L/100$, its edge-connected block image would therefore have diameter at least

$$M/100 - 1 \geq M/200$$

for sufficiently large M. That image is contained in the visited set in its slab, contradicting the applicable V_j. The whole-slab definition covers both configurations immediately before and after every physical jump. Almost surely no ring occurs at a deterministic slab boundary. At time zero the syndrome is empty. These observations give the assertion for all times in the stated interval. □

Proposition 5.3 (Cumulative finite-torus estimate)**.** *There exist $\eta_0, C_0, \beta > 0$ and an integer L_0, independent of the allowed spatial noise rates, such that in either sector, for every $L \geq L_0$ and deterministic $t \geq 0$,*

$$\mathbb{P}\{\exists s \in [0,t] : \textit{a syndrome component at } s \textit{ has diameter } \geq L/100\} \leq \min\{1, C_0(1+t)e^{-\beta L}\}. \tag{5.6}$$

The initial syndrome is empty.

Proof. For $t > 0$ put $n = \lceil t/(hT) \rceil$. A union bound over the events in Lemma 5.2, with Lemma 5.1, gives

$$CnM^4\big[e^{-\kappa h} + 2he^{-\kappa M/200}\big] \tag{5.7}$$

as an upper bound on the probability. Some of these tests extend past t. Using their future marks in an upper bound is legitimate; neither the device nor the recovery consults them.

For all sufficiently large M, $h \geq M/[200(R+3)]$, while $h \leq M$. Also $M \geq L/(2B)$ once L is sufficiently large. Consequently the bracket in (5.7) is bounded by a polynomial in L times $e^{-b_1 L}$ for a fixed $b_1 > 0$. Since $hT \geq 1$, we have $n \leq 1 + t$. Absorb all remaining polynomial factors into half the positive exponential rate, enlarging L_0 and C_0. This gives (5.6) with, for example, some $\beta \leq b_1/2$. The event at $t = 0$ is empty. The probability is always at most one.

Overlapping windows share many increments. Their independence was not used: the entire accumulation step is a union bound on the original sample of marks. The argument also never assumes contraction for arbitrary comparison data at a restart. □

6 A fixed current-syndrome decoder

We construct a deterministic function $D(s)$ with $\partial_2 D(s) = s$ on every attainable syndrome s. The same definition is used on the primal and dual complexes, with separate fixed orderings of their cells. No mark times enter this definition.

6.1 Filling a small component

Set $\delta = L/100$. Fix representatives of torus vertices and a lexicographic ordering of all torus plaquettes. Suppose every edge-connected component of s has ambient torus diameter less than δ. For a nonempty component Γ, choose its lexicographically least vertex v_Γ. Lift every vertex by its shortest coordinate displacement in $(-L/2, L/2]$ from this anchor, and let K_Γ be the coordinate box between the lifted coordinate extrema.

Each lifted vertex lies within δ of the anchor. The box has each side length at most 2δ; its projection embeds for large L. Neighbouring vertices lift consistently: a coordinate difference between their lifts is at most $2\delta < L - 1$, which rules out the spurious displacement $L - 1$ across a unit torus edge. All edges of Γ therefore lift to edges of this box. Even degree at every vertex makes the component a cellular 1-cycle.

Here is an explicit reason that this cycle fills. A cubical box is a product of integer intervals, allowing intervals of length zero. Collapse the final layer in one nondegenerate coordinate interval onto the preceding layer. On cellular chains let p be that projection, and let κ_h send a horizontal cell in the final layer to its prism across the collapsed unit interval; let κ_h vanish on other cells. Directly on each cell,

$$\partial\kappa_h + \kappa_h\partial = \mathrm{id} + p$$

over $\mathbb{F}_2$: the two horizontal faces of a prism give the cell and its projected copy, while the side-face terms cancel. For a cell crossing the layer, its top horizontal face supplies that prism in $\kappa_h\partial$, giving the same identity. Iterating the collapses in all coordinates contracts the box to a vertex. Thus every positive-dimensional cycle in the box is a boundary, including every 1-cycle and 2-cycle. This also covers degenerate boxes.

Choose the lexicographically least plaquette subset in K_Γ whose boundary is Γ, with ordering induced by its projection to torus cells. A finite set of fillings is nonempty by the contraction, so this choice is well defined. Project back to the torus and add the component fillings modulo two. Call the resulting chain $D(s)$. Set $D(0) = 0$. If some component has diameter at least δ, use instead the lexicographically least global solution of $\partial_2 D = s$. Attainability means $s \in \operatorname{im}\partial_2$, so a solution exists in that case as well.

The construction is fixed once the cell orderings are fixed. In particular an unchanged small component has exactly the same filling at every observation. This consistency, rather than optimal area, is its essential property.

6.2 Invariance across a physical face flip

For a current error chain E, define its filled homology class by

$$\mathfrak{h}(E) = [E + D(\partial_2 E)] \in H_2(K_L; \mathbb{F}_2). \tag{6.1}$$

The bracket is meaningful since the displayed chain is closed.

Lemma 6.1 (Local preservation of filled homology)**.** *For all sufficiently large L, let $E' = E + p$ differ from E by one plaquette. If every syndrome component of both $s = \partial_2 E$ and $s' = \partial_2 E'$ has diameter less than δ, then*

$$\mathfrak{h}(E') = \mathfrak{h}(E).$$

This holds for arbitrary merges, splits, creation or disappearance of components at that flip, and on the dual complex as well.

Proof. The only modified syndrome edges lie in ∂p. Any component in either decomposition that does not meet a vertex of p remains unchanged. Indeed its edges and their incident connections are untouched; a changed connection would have an endpoint on p. Conversely, a component changed by deletion or insertion has a path to a modified edge and hence meets a vertex of p. The fillings of all unchanged components cancel in $D(s) + D(s')$.

Lift p once to the integer lattice. Align the anchor lift of each affected component using one of its vertices on this face. Its anchor is within δ of that shared vertex, and its filling box is within δ in each coordinate of the anchor. Thus all affected fillings, together with p, are supported in one integer box within coordinate distance $2\delta + 2$ of a fixed vertex of p. For large L the width $4\delta + 4$ is less than $L - 1$, so the common box embeds in the torus.

The chain

$$C = p + D(s) + D(s')$$

has boundary $\partial p + s + s' = 0$. Its uncancelled support is in the common box. The contraction above makes C the boundary of a 3-chain there. Consequently adding C does not change the homology class in (6.1). Duality preserves all these incidence and support properties, proving the last assertion. □

Corollary 6.2 (Retention along the marked history)**.** *In either sector, started with $E = 0$, the probability that $\mathfrak{h}(E_s) \neq 0$ at any $s \leq t$ is at most $C_0(1+t)e^{-\beta L}$, with constants as in Proposition 5.3. An initial closed chain E_0 instead gives retention of its initial filled class with the same bound.*

Proof. On a finite torus there are finitely many clock rings up to each finite t, almost surely, even if that deterministic t is exponentially large in L. Order all effective face flips. If the large-component event of Proposition 5.3 does not occur, each flip has only small components on both sides. Lemma 6.1 propagates the initial class through all of them. With $E = 0$ it starts at zero. For an added closed initial chain, the syndrome and the enabled correction predicates are identical under common marks; the decoded class is shifted by its fixed initial class. The same large-component event therefore controls a change in that class. □

The decoder need not infer an initial logical class from the syndrome. An encoded start supplies that class. More generally, a syndrome-only recovery can preserve a known encoded input through this dynamics but cannot distinguish two initial errors differing by a hidden logical membrane.

7 Quantum evolution and worst-case fidelity

We verify directly that the two sector histories analyze (2.4). This also checks the continuous-time CSS hypotheses in the forward memory implication of Drechsler–Oreg–Aharonov, Theorem 1 (Section 1.5.1; formal version Section 3.1) [8], without requiring that implication as a probabilistic input.

Lemma 7.1 (Pauli-trajectory realization)**.** *For every finite L, every encoded density operator ρ, possibly entangled with an arbitrary reference, and every $t \geq 0$, the state $(e^{t\mathcal{L}_L} \otimes \mathrm{id})(\rho)$ is the mixture over the independent marked histories specified in Section 2, applying their cumulative physical Paulis to ρ. Each sector has the single-face classical correction rule of Section 3 and noise rate at most 2η. The probabilities of these histories are independent of ρ.*

Proof. On a Pauli translate $U\rho U^\dagger$ of any encoded state, every stabilizer has a definite sign determined by U, regardless of the logical state or the reference. Thus any tested projector Π has a definite scalar value $b \in \{0, 1\}$ on this whole branch. For $J = P\Pi$, $J^\dagger J = \Pi$, and

$$\mathcal{D}[P\Pi](U\rho U^\dagger) = \begin{cases} 0, & b = 0, \\ PU\rho U^\dagger P - U\rho U^\dagger, & b = 1. \end{cases}$$

A rate-one attempt clock gives exactly this infinitesimal transition law. A noise clock with rate $\gamma_{p,P}$ gives $\gamma_{p,P}(P_p \rho P_p - \rho)$. Summing these transition laws over the finitely many clocks shows that the marked-history mixture solves the Lindblad equation with the prescribed initial state. Finite-dimensional uniqueness of this linear differential equation identifies it with the quantum evolution. This uniqueness follows, for example, by the convergent matrix exponential of the finite-dimensional superoperator.

Products of physical Paulis stay physical Paulis up to a phase. Their syndrome signs depend only on the error chains, so neither the correction rates nor the history probabilities depend on the logical state. A primal correction adds its plaquette to the primal error chain, and a dual correction adds its dual plaquette to the dual error chain. The noise marginals are the sums of the Z, Y and X, Y rates respectively. Shared Y marks can correlate the sectors, but do not alter either marginal law. □

The stipulated generator is a correction Lindbladian in [8, Definition 2.24]: its Hamiltonian vanishes, its Z jumps use only X-check functions, and its X jumps use only Z-check functions. Its noise has Definition 2.25's independent Pauli form. The initial encoding has all check signs positive, and D is a compatible decoder with $D(0) = 0$ as in their Definition 2.5. These checks explain the applicability of their forward Theorem 1 (Section 1.5.1; formal version Section 3.1). The direct realization just proved also handles the stronger reference statement used here.

To implement $\mathcal{R}_L$, measure the commuting checks, apply $Z^{D_{\mathrm{primal}}(s_X)}$ and $X^{D_{\mathrm{dual}}(s_Z)}$, invert the code encoding on the corrected code space, and output the designated qubit. The order of the two Pauli corrections changes only a phase. Each correction removes its associated syndrome and

commutes with the other family of checks, so their combination returns every attainable branch to the code space.

Proof of Theorem 2.1. Apply Corollary 6.2 to the two marginal sectors in Lemma 7.1. A union bound, requiring no independence between sectors, gives an event $\mathcal{S}_t$ with

$$\mathbb{P}(\mathcal{S}_t) \geq 1 - 2C_0(1+t)e^{-\beta L}$$

on which both filled classes remain zero through t. On such a history the primal and dual residual error chains are boundaries of 3-chains. Their Paulis are products of Z and X stabilizers, respectively. The total residual acts as the identity on every code vector, up to a physically irrelevant phase. It therefore restores all six logical qubits exactly, including their joint state with a reference.

Since the history probabilities do not depend on the input, the effective channel is a convex mixture whose contribution from $\mathcal{S}_t$ is exactly $\mathbb{P}(\mathcal{S}_t)\,\mathrm{id}$. For any pure $|\Psi\rangle_{QR}$ the target overlap on the other branches is nonnegative. Hence

$$\langle\Psi|(\mathcal{N}_{L,t} \otimes \mathrm{id}_R)(|\Psi\rangle\langle\Psi|)|\Psi\rangle \geq 1 - 2C_0(1+t)e^{-\beta L}. \tag{7.1}$$

This conclusion is uniform in the fixed unused logical state σ; on success that state too is restored before it is discarded.

Choose $c = \beta/3$. For $0 \leq t \leq e^{cL}$,

$$2C_0(1+t)e^{-\beta L} \leq 4C_0 e^{-(\beta-c)L} \leq 4C_0 e^{-cL}.$$

Take $C = 4C_0$, enlarging it if needed for a common L_0. With a trivial reference, minimizing (7.1) over all pure inputs yields (2.6). With a general reference it gives the asserted stronger inequality. All geometric and probabilistic constants were fixed independently of L, the rate values within the bound, and the input. The fillings use only the syndrome and fixed cell orderings. This proves every claimed uniformity. □

For a Bell input, the same argument gives entanglement fidelity at least $1 - Ce^{-cL}$. The maximum overlap of a pure two-qubit product with a Bell vector is $1/2$, by the Cauchy–Schwarz inequality applied to its two Schmidt terms. Convexity extends this upper bound to separable states. Thus for sufficiently large L the guaranteed Bell fidelity is above that classical benchmark throughout the interval. The worst-case statement in Theorem 2.1 does not rely on choosing this particular input.

8 The process on the infinite lattice

We now give the construction and prove Theorem 2.2. No global chain of all events is ordered on the infinite lattice; there are infinitely many marks in any positive time interval. Local chronological dependency estimates provide the appropriate construction.

8.1 Graphical existence

Lemma 8.1 (Finite influence on bounded intervals)**.** *For independent rate-one correction clocks and independent sector noise clocks of rate at most 2η on each plaquette of $\mathbb{Z}^4$, finite-volume constructions started empty agree on any fixed finite target set, throughout each bounded time interval, for all sufficiently large cuts, almost surely. Their common limit is a local right-continuous process with left limits. It preserves even degree at every syndrome vertex.*

Proof. Superpose the correction and noise marks of a face. Its rate is at most $\Lambda = 1 + 2\eta$, uniformly over faces. Fix a finite target set F and time length $H_* > 0$. If a discrepancy at a remote boundary affects an edge of F, tracing dependencies backwards through the local maps produces a chronological chain of intersecting marked faces ending next to F. A correction reads only its two tested edges and toggles its boundary, so any discrepancy it creates has such a predecessor. A noise mark acts by the same fixed toggle in both constructions and cannot create a discrepancy from equal inputs. Allowing it as a possible dependency step only enlarges the set of chains being counted.

There are at most C_F terminal faces next to F, and at most 96 predecessor face labels for each subsequent step. For any specified sequence of m labels, the ordered factorial-moment integral is bounded by $(\Lambda H_*)^m/m!$, including repeated labels. Therefore

$$\mathbb{P}\{\text{such a chain of length } m\} \le C_F \frac{(96\Lambda H_*)^m}{m!}. \tag{8.1}$$

The series of the right sides over m converges. A dependency from distance tending to infinity requires length tending to infinity, because consecutive face anchors differ by at most one in each coordinate. By the summability in (8.1), almost surely sufficiently long chains do not occur. Hence all sufficiently large finite cuts agree on F throughout $[0, H_*]$. Any boundary convention confined to the cut boundary has the same limit.

Intersect these probability-one events over increasing finite target boxes and integer H_*. The resulting limits agree on overlaps and define consistent local paths for all $t \ge 0$. Each fixed edge is incident to finitely many face processes and is therefore touched only finitely often on a bounded interval. The paths are right-continuous with left limits. Every vertex initially has degree zero, and each local face toggle changes its degree by zero modulo two. This parity property passes from the finite constructions to the local limit. □

The same construction can be made simultaneously in the two marginal sectors with shared Y marks. The assertions here need only the construction and law of one named sector. They do not require a definition of an infinite-volume encoded quantum state.

8.2 Uniform component tails

Use ordinary blocks

$$Q_z = Bz + \{0, \dots, B-1\}^4, \qquad z \in \mathbb{Z}^4,$$

with the fixed B, T, R already chosen for a sufficiently small noise threshold. The local tests and clean cuts of Section 3 are finite objects, so they apply to the graphical limit. Their probabilities and predecessor implications are unchanged. By induction, physical visits satisfy (4.3) for every finite slab index.

Every comparison root at finite date has a finite dependency cone. There is no periodic identification and no bound on its permissible depth. Thus Proposition 4.1 gives the certificates uniformly in the root date. Proposition 4.2 applies on the entire half-space of dates $k > 0$. One can also obtain its half-space assertion by applying the finite-path argument directly: each finite witnessing path has a finite maximum date, and the deterministic chain count already includes all possible center dates with the stated distance restrictions.

Proof of Theorem 2.2. Existence and parity follow from Lemma 8.1. Fix x, t, d, and let z_x be the block containing x. Take the slab containing t, with index k, as specified after the theorem. Suppose

at a common instant of that slab x belongs to a component of diameter at least d. For a finite component, the triangle inequality implies that some vertex has distance at least $d/2$ from x; if the diameter is only a supremum one can decrease this target radius arbitrarily and pass to a limit of the corresponding events. An unbounded component likewise supplies a finite connecting edge path to any fixed radius.

The blocks along a connecting physical edge path form a spatial path in $P^0(\cdot, k)$ by (4.3). Starting at z_x, it reaches coarse distance at least $d/(2B) - 1$. Proposition 4.2 bounds the probability by

$$C \exp\left[-\kappa \max\{0, d/(2B) - 1\}\right] \leq C' e^{-\mu d}$$

for $\mu = \kappa/(2B)$, after enlarging C'. This proves the whole-slab estimate, uniformly in x and k. The event at the specified time t is a subset. At $t = 0$ the instantaneous syndrome is empty; the whole-slab estimate uses the slab beginning there. These statements prove (2.7) with uniform constants.

For the low-noise density limit it is useful that the cost A in Proposition 4.1 is adjustable. At any specified block and date,

$$\mathbb{P}(P^0(z, k)) \leq \sum_{r \geq 1} e^{-Ar} = \frac{1}{e^A - 1}.$$

Given $\epsilon > 0$, choose A large enough to make this quantity less than ϵ, and large enough for the path estimate if desired. Next choose q, then B, and then a positive threshold η_ϵ by the certificate proof and Lemma 3.5. For every family of rates bounded by η_ϵ, occupation of x at any deterministic time implies this block proxy event in its slab. Hence its probability is at most ϵ, uniformly in x, t and the rate family. This is precisely the stated limit. The block scale in this density argument may depend on ϵ; the argument asserts no uniform limiting value for the decay constant μ as the threshold is varied. □

For each specified vertex and time the diameter tail also rules out membership in an infinite component by letting d increase. This is a rooted estimate and carries no factor for the volume of the lattice. At a fixed positive noise rate new local syndromes continue to be created, so uniform-in-time tails do not mean that the entire infinite syndrome eventually becomes empty.

9 Spacetime local stochastic Pauli noise

This section extends the finite-torus memory theorem. The correction rule, correction-clock marginal and decoder are unchanged. The noise is now specified by a classical joint trajectory law rather than, in general, by an additive Lindblad noise generator.

9.1 Lindbladian correction, trajectory-defined noise

Fix integers $g \geq 1$ and $r_0 \geq 0$, independently of L. There are at most g anchor labels at each spatial lattice vertex. Each anchor α has a fixed footprint of data qubits within sup radius r_0 of its position. Enlarge this fixed radius if necessary to contain the vertices of both the associated primal and dual faces. The time unit is the inverse correction rate, as in Section 2. Elementary locations are

$$(\alpha, j), \qquad \mathsf{I}_j = (j, j+1], \qquad j = 0, 1, \ldots.$$

Let $\mathcal{F}$ be a random set of such locations. For a strength $0 < p < 1$, impose the spacetime local-stochastic condition

$$\mathbb{P}\{S \subseteq \mathcal{F}\} \le p^{|S|} \quad \text{for every deterministic finite set } S \text{ of distinct elementary locations.} \tag{9.1}$$

An error at a time in I_j is a Pauli string whose entire support is contained in the union of footprints of the faulty locations $(\alpha, j) \in \mathcal{F}$. Its labels, times and multiplicities may otherwise be arbitrary, subject to local finiteness on bounded time intervals. A single faulty location may apply several nearby single-qubit factors at once, paying for that location once. Even an extensive simultaneous string is allowed when it has the required cover by faulty locations in its own time cell.

The two correction-clock families retain their joint product law of independent rate-one Poisson processes. Noise may correlate across positions, dates, Pauli types and sectors, and with the correction clocks; only this correction-clock *marginal* is fixed. The complete classical noise law must be independent of the hidden logical input and its reference. Adaptive choices based on exogenous randomness, the syndrome and previous Pauli-and-clock history are allowed when their joint fault set obeys (9.1). No extra measurement of the logical subsystem is part of a noise history.

A definite order at ties between faults and correction attempts is part of the history. For proofs, factor each simultaneous Pauli string into its nonidentity single-qubit factors in a fixed internal order, inserting no correction attempt between them. All these virtual stages have the original time and fault cover. Readout at a deterministic time is after all events at that time. The encoded state at zero precedes every allowed error. These conventions allow noise pulses at deterministic time-cell boundaries.

We call this model *Lindbladian correction, trajectory-defined noise*. Each correction attempt is the Poisson unravelling of $\sum_p (\mathcal{D}[Z_p \Pi_p^X] + \mathcal{D}[X_p \Pi_p^Z])$ on a syndrome-definite branch. A general correlated history need not yield a Markov noise process or a semigroup for the combined evolution. Its ensemble is defined directly by averaging its Pauli trajectories on encoded inputs. The input-independence hypothesis is essential to that channel interpretation.

Lemma 9.1 (Dependence on the time grid)**.** *For any law satisfying* (9.1) *on unit cells with parameter p, a unit grid shifted by any fixed amount admits a covering fault set with local-stochastic parameter at most $2\sqrt{p}$. Aligned subdivision of each unit cell into k equal cells, for any integer $k \ge 1$, admits parameter at most $p^{1/k}$. The spatial footprints can be kept unchanged. Bounds exceeding one are understood as vacuous.*

Proof. On the shifted grid declare a location faulty if an old faulty cell with the same anchor intersects its interval. At time zero, intersect intervals with the nonnegative time axis; old cells before zero can be treated as empty. Each new interval intersects at most two old intervals, and each old interval intersects at most two new ones, using the same half-open endpoint convention. For a prescribed set of m new locations, expand at most 2^m choices of an old covering key per location. Every such choice contains at least $m/2$ distinct old keys. Its probability is at most $p^{m/2}$ by (9.1). A union bound gives $2^m p^{m/2} = (2\sqrt{p})^m$. The support cover is preserved because the old cell containing an error time intersects the new cell containing that time.

For aligned refinement, declare all k children of a faulty old location faulty. A prescribed set of m refined locations has at least m/k distinct old parents, so its joint probability is at most $p^{m/k}$. The same footprints cover every original error. These are changes of the strength parameter, not merely a fixed rescaling of constants. □

Theorem 9.2 (Exponential lifetime under spacetime local stochastic Pauli noise)**.** *For every fixed g, r_0 there exist $p_*, c, C > 0$ and an integer L_0, independent of the admissible joint history law and of the logical input, such that the following holds. For every integer $L \geq L_0$, every history law with the correction-clock marginal, bounded footprints, local finiteness and input-independence specified above, and strength $p \leq p_*$ in* (9.1) *on unit time cells, the current-syndrome decoder of Section 6, from an encoded start, obeys*

$$\inf_{0 \leq t \leq e^{cL}} \min_{\|\psi\|=1} \langle \psi | (\mathcal{R}_L \circ \mathsf{Store}_t \circ \mathcal{E}_{\mathrm{enc}})(|\psi\rangle\langle\psi|) |\psi\rangle \geq 1 - Ce^{-cL}. \tag{9.2}$$

Here Store_t is the trajectory-defined storage ensemble. The same bound holds for overlap with any pure input and untouched reference, uniformly in the designated logical qubit and fixed state of the other five. It allows bounded-range multi-qubit and adaptive faults satisfying the stated joint condition, and concerns stochastic Pauli noise only.

The proof occupies the remaining subsections. All estimates for correlated faults are unconditional at deterministic dates; a conditional bound after an arbitrary revealed past is not assumed.

9.2 Noise models included in the hypothesis

Independent single-qubit Pauli Poisson jumps of deterministic rates $\gamma_{p,P}(t) \leq \eta$ satisfy (9.1) with $p \leq 1 - e^{-3\eta} \leq 3\eta$: use one anchor per qubit and mark its unit cell when any noise jump occurs. Distinct qubit-time cells are independent by the product Poisson law. Thus the extension includes the constant-rate model of Theorem 2.1, and its uniformly bounded deterministic time-dependent version. Fresh product Pauli layers with each one-qubit nonidentity probability at most 3η, with at most n_l prescribed layers per unit interval, likewise give $p \leq 3n_l\eta$ by a cell union bound and independence between distinct cells. Here n_l is fixed independently of L.

Independent bounded-range multi-qubit Pauli jumps are another special case. Assign each jump type to a fixed anchor in its support, with its support inside that anchor's footprint. Suppose the sum of deterministic intensities at each anchor is at most λ. The event that any such jump occurs in a unit anchor cell has probability at most $1 - e^{-\lambda}$; these events are independent across anchors and unit intervals. This proves (9.1) with $p \leq 1 - e^{-\lambda}$. A bound λ on the sum of jump rates touching each qubit also suffices when qubits are the anchors, since the types assigned to a qubit are a subset of the types touching it. For independent constant-rate jumps added to the correction, the usual generator is the correction Lindbladian plus $\sum_\tau \lambda_\tau \mathcal{D}[P_\tau]$. The Pauli-branch calculation in Section 7 proves this realization directly on encoded inputs. The general correlated model makes no such sum-generator assertion.

9.3 Geometry and a joint coarse-bin cost

Keep $R = 2000100$. Take an integer $B \geq \max\{10, r_0 + 10\}$, and use M, b, T, m_f, Q_z, U_z from (3.3), with $T = 1000B$ an integer. We require L sufficiently large for the radius-$(R+6)$ block neighbourhoods to lift injectively. In the dual sector use its fixed half-cell coordinate translation.

Let $\mathcal{C}(z, k)$ be failure of either of the two clock-only tests in Definition 3.4: the exclusion of a chronological chain of m_f attempts and completion of every target clearance schedule by T. Its probability is bounded using only the correction-clock marginal by

$$\pi_B := CB^4(96e/1000)^{1000T} + (2B+1)^4(4 \cdot 2^{12})^{8B+2} e^{-T/2} \longrightarrow 0. \tag{9.3}$$

This is exactly the clock part of the proof of Lemma 3.5: ordered repeated labels use Poisson factorial moments, and the schedule uses distinct-vertex exponential waits. Correlation with faults does not alter this marginal estimate.

Let $\mathcal{V}(z,k)$ be the deterministic set of elementary anchor-time cells in slab I_k whose possible footprints can include a face incident on U_z. Use the full allowed radius-r_0 collar, rather than the realized fault support. Define the new bad-bin indicator

$$J_{\mathrm{ls}}(z,k) = \mathbf{1}\{\mathcal{C}(z,k) \text{ or } \mathcal{F} \cap \mathcal{V}(z,k) \neq \varnothing\}. \tag{9.4}$$

The elementary time cells partition the slab because T is integer. Put

$$H_0 = R + 6, \qquad D_0 = (4H_0+1)^4, \qquad V_B = gT\left\lceil 2B(2R+1) + 2r_0 + 12\right\rceil^4. \tag{9.5}$$

All clock inputs and spatial anchors of $\mathcal{V}(z,k)$ are within H_0 coarse blocks of z, and $|\mathcal{V}(z,k)| \le V_B$. Indeed U_z has coordinate width at most $2B(2R+1)$; the footprint collar and face-vertex allowance add $2r_0+12$, including the fixed dual translation. Since $B \ge r_0 + 10$, these physical collars fit in the stated fixed coarse margin. A larger fixed cell-coordinate convention can be absorbed by increasing that margin and H_0 once.

If $J_{\mathrm{ls}}(z,k) = 0$, there is no noise face incident on U_z, even at an intermediate factor of an atomic Pauli. The backward discrepancy argument in Lemma 3.6 consequently gives agreement with the clean cut throughout the slab. A discrepancy from its boundary would need a chronological chain from distance at least $RB - 2 > m_f + 5$, contradicting the clock test. The coordinate confinement and clearance proofs then give all five predecessor implications of that lemma for every even starting syndrome. They also give its whole-slab visit implication, now including the virtual intermediate states. A distant atomic string changes no edge in U_z when the test is good. At a slab boundary the starting state for the next slab is after all marks of the preceding slab, so deterministic fault-time pulses cause no exception.

Lemma 9.3 (Joint bad-bin bound)**.** *For every $b_* \in (0,1)$, one can choose B as above and then $p_* > 0$, depending only on g, r_0, b_*, so that for every admissible law at $p \le p_*$ and every deterministic finite set $\mathcal{U}$ of distinct coarse spacetime bins,*

$$\mathbb{P}\{J_{\mathrm{ls}}(u) = 1 \text{ for all } u \in \mathcal{U}\} \le b_*^{|\mathcal{U}|}. \tag{9.6}$$

This assertion is uniform over the dates in $\mathcal{U}$.

Proof. Write $m = |\mathcal{U}|$, with the empty case immediate. Join bins with the same slab index and coarse spatial distance at most $2H_0$. A greedy selection retains $\ell \ge m/D_0$ nonadjacent bins. Their potential elementary fault sets $\mathcal{V}$ are disjoint, and their clock-only tests use disjoint Poisson increments. The latter tests are independent under the product clock marginal, regardless of their correlations with the noise.

Partition these ℓ bins into a specified set $\mathcal{A}$ requiring clock failure and a set $\mathcal{B}$ requiring a fault-envelope hit. The clock intersection has probability at most $\pi_B^{|\mathcal{A}|}$. For the other intersection choose one elementary key from each $\mathcal{V}$. Those keys are distinct because the envelopes are disjoint. Equation (9.1) and a union bound over at most $V_B^{|\mathcal{B}|}$ choices give an upper bound $(V_B p)^{|\mathcal{B}|}$. Cauchy–Schwarz for the two event indicators bounds their intersection by $\pi_B^{|\mathcal{A}|/2}(V_B p)^{|\mathcal{B}|/2}$, without conditioning either marginal on the other.

Sum over the partitions. Every selected bin is bad only if it has at least one of these two causes. Therefore the probability that all original bins are bad is at most

$$(\sqrt{\pi_B} + \sqrt{V_B p})^{\ell}.$$

Choose B so large that $\sqrt{\pi_B} \le b_*^{D_0}/2$, using (9.3), and then $p_* > 0$ so small that $\sqrt{V_B p_*} \le b_*^{D_0}/2$. The displayed expression is at most $b_*^{D_0 \ell} \le b_*^m$. Only the unconditional fault inclusion bound and the clock marginal were used. □

9.4 Joint costs for localized witness families

Use (4.1) and (4.2) with J_{ls} in place of J, writing the resulting fields as H^a_{ls} and P^a_{ls}. All-one initialization at the deterministic date a is unchanged. The deterministic implications just proved give, from the encoded start,

$$\begin{aligned} Q_z \text{ occupied at } kT &\implies H^0_{\mathrm{ls}}(z,k) = 0, \\ Q_z \text{ visited in slab } k &\implies P^0_{\mathrm{ls}}(z,k). \end{aligned} \tag{9.7}$$

The visits in the second line include every single-face intermediate stage of each simultaneous Pauli string. The induction is the one in (4.3): endpoint occupation in a good bin has a zero predecessor in every A_s, and any visit has at least one in Δ.

Proposition 9.4 (Joint localized-witness cost)**.** *For every $A > 0$, the parameters can be chosen so that the following holds in every deterministic window of depth w after a date a, with $2Rw < M/2$. At each site u of that window there are certificate events $\mathcal{E}^a_{\mathrm{ls}}(u,r)$, with allowed radii*

$$r = K_0(n+1), \qquad K_0 = 10(R+6), \qquad n = 0, 1, \ldots,$$

such that $P^a_{\mathrm{ls}}(u)$ implies at least one event. For every finite predetermined family of mutually disjoint closed ρ-balls $\mathcal{B}_\rho(u_j, r_j)$, with centers in that window and allowed radii, their joint cost satisfies

$$\mathbb{P}\left(\bigcap_j \mathcal{E}^a_{\mathrm{ls}}(u_j, r_j)\right) \le \exp\left(-A \sum_j r_j\right). \tag{9.8}$$

The events are unions of requirements on bad sink bins localized inside their named balls. No independence between the events is asserted.

Proof. A bad bin at u gives a one-sink candidate with $n = 0$. Otherwise $P^a_{\mathrm{ls}}(u)$ supplies a zero predecessor $(y, k-1)$ within spatial distance R. There is no zero at the all-one bottom date. At a later date lift that predecessor's backward cone to the integer lattice, copy its maps, and put constant-one maps at the bottom and everywhere outside the cone. The recursion is finite and time-directed, so its root value is the unique one from its bottom data and hence the maximal trajectory value.

This graph satisfies all the hypotheses of the deterministic SST Theorem 22 checked in Section 4: each nonconstant map has finite nonempty type-s predecessor set A_s, constants have no outgoing dependencies, and each dependence edge decreases time by one. That theorem supplies a connected contour at a zero. SST Definitions 17 and 21 and Eq. (2.16) make its sinks distinct bad sites and exclude constant-one cutoff sites [15]. At each date the cone width is at most $2Rw < M/2$, so

distinct lifted sinks at that date have distinct physical bin indices. Distinct dates already give distinct indices. This argument is deterministic and has no fault-independence hypothesis.

Let n be the number of edges per charge and m the number of sinks. The linear forms (4.6) have increments at least $0, 0, 0, 0, 1$ on the five prescribed types, sum to zero, and have respective losses at most $R, R, R, R, 4R$ on arbitrary offsets in Δ. Consequently the flux calculation (4.7), using SST Lemmas 38 and 40, gives $m \geq 1 + n/\alpha$, where $\alpha = 1 + 8R$. The enumeration (4.8) uses SST arXiv:2202.10999v3, proof Eq. (5.28); its weaker Lemma 39 statement Eq. (5.15) would also suffice. After the bounded predecessor choice and the direct case, the number of candidates at u is at most $C_0 C_1^n$, with constants depending only on R.

Define $\mathcal{E}^a_{\mathrm{ls}}(u, K_0(n+1))$ by taking the union over those deterministic cone candidates and requiring only that each named sink is bad. A connected contour has $5n$ edges. Hence every sink is within $5Rn + R$ spatial units and $5n + 1$ dates of u, and lies in its named ball. Even the local clock/fault inputs of its bad test fit after adding the H_0 collar. Empty events can be used at other radii. The contour implication gives the claimed cover.

For the prescribed disjoint balls, expand the intersection of the candidate unions into tuples of deterministic contours. Within one candidate the sinks are distinct by the embedding and lift argument. Across candidates in a tuple they are distinct because the balls are disjoint as sets on the actual torus times time. Apply Lemma 9.3 once to the *union* of all sink bins in that tuple. With $n_j = r_j/K_0 - 1$, summing over the tuples gives

$$\mathbb{P}\left(\bigcap_j \mathcal{E}^a_{\mathrm{ls}}(u_j, r_j)\right) \leq \prod_j \left[C_0 C_1^{n_j} b_*^{1+n_j/\alpha}\right]. \tag{9.9}$$

The displayed product is a counting bound: the number of choices multiplies and the number of distinct sinks adds. It does not factor the underlying noise law.

Choose $b_* > 0$ sufficiently small so that $C_0 C_1^n b_*^{1+n/\alpha} \leq e^{-AK_0(n+1)}$ for every $n \geq 0$, by bounding the constant term and the coefficient of n in its logarithm. Then choose B, p_* by Lemma 9.3. This proves (9.8). Balls of radius above $M/2$ may wrap spatially around the torus. The sink sets are physical indices without repeated copies, the quotient volume bound remains valid, and only disjoint actual balls are used. No independence claim is needed for wrapping or nonwrapping balls. □

Proposition 9.5 (Centered crossing under correlated faults)**.** *For sufficiently large fixed A, and the resulting choices of parameters, there are $C, \kappa > 0$ such that for every deterministic window above, every specified anchor u in it, and every real $s \geq 0$,*

$$\mathbb{P}\left\{\begin{array}{l} \text{a finite path in } P^a_{\mathrm{ls}} \text{ starts at } u, \text{ stays in the window,} \\ \text{and reaches } \rho\text{-distance at least } s \end{array}\right\} \leq Ce^{-\kappa s}. \tag{9.10}$$

Path steps have ρ-distance at most R, and the constants are uniform over dates, L, and admissible laws.

Proof. For a realizing finite path, assign an occurring certificate ball to each distinct site. Greedily select disjoint balls in decreasing radius. A discarded ball meets a selected one with at least its radius, so the selected triple enlargements cover the path. The covers of its endpoints are connected by a simple chain of selected centers whose successive distances are at most $3r + 3r' + R$. If v is the sum of radii in this chain and j its length, endpoint displacements and the triangle inequality give

$s \le 6v + R(j-1) \le 7v$, since every radius exceeds R. This is a product-metric argument, including on the torus.

For fixed ordered radii, the initial center has at most $(6r_1+1)^5$ choices around the specified anchor, and subsequent centers have at most $(6r_i + 6r_{i+1} + 2R + 1)^5$ choices. As in the explicit entropy calculation in Proposition 4.2, their product is at most e^{60v}. There are at most 2^v ordered compositions of v. For each deterministic chain with disjoint original balls, use the joint estimate (9.8) to pay e^{-Av} for all its certificates at once. Thus, taking $A > 100$,

$$\mathbb{P}(\text{crossing}) \le \sum_{v \ge \max\{1, \lceil s/7 \rceil\}} e^{-(A-60-\log 2)v} \le Ce^{-\kappa s}.$$

Adaptive selection after observing the path causes no probability factorization: the union bound enumerates deterministic chains before applying their joint cost. □

9.5 Unconditional overlapping windows

Proposition 9.6 (Correlated cumulative history bound)**.** *For each named CSS sector there exist $C_0, \beta > 0$ and L_0, uniform over the admissible history laws at $p \le p_*$, such that from an encoded start, for every $L \ge L_0$ and deterministic $t \ge 0$,*

$$\mathbb{P}\left\{ \begin{array}{l} \text{a sector syndrome component has diameter } \ge L/100 \\ \text{by time } t, \text{ including virtual single-face stages} \end{array} \right\} \le \min\{1, C_0(1+t)e^{-\beta L}\}. \quad (9.11)$$

Diameter is ambient torus sup diameter.

Proof. Set $h = \lfloor M/[100(R+6)] \rfloor$ and $a_j = jh$. For all sufficiently large M, $h \ge 1$ and a two-epoch window of depth $2h$ satisfies $4Rh < M/2$. On the original joint marks define $H_j = H_{\mathrm{ls}}^{a_j}$, all ones at its own deterministic start. Define G_j as equality of H_j, H_{j+1} everywhere at a_{j+2}, and V_j as absence of a spatially connected set of diameter at least $M/200$ in $P_{\mathrm{ls}}^{a_j}(\cdot, k)$ at any of the $2h$ window dates.

During the overlap $H_j \le H_{j+1}$, by all-one initialization of the latter and monotonicity of each common update. A disagreeing endpoint has a disagreeing predecessor unless its map is constant, which cannot transmit disagreement. Trace one such predecessor per date down to a_{j+1}. The resulting path consists of zeros of the older comparison, hence sites in $P_{\mathrm{ls}}^{a_j}$; it has temporal span h and steps at most R. Proposition 9.5 and a union over final sites give

$$\mathbb{P}(G_j^c) \le CM^4 e^{-\kappa h}.$$

If V_j fails, a spatial path in one of those dates reaches distance $M/200$ from some initial block. Summing the same path bound over dates and anchors gives

$$\mathbb{P}(V_j^c) \le 2ChM^4 e^{-\kappa M/200}.$$

These are unconditional bounds at each deterministic j. The joint inclusion bound applies to any specified dates, so no fresh conditional noise law at a_j is needed.

On $\bigcap_{0 \le j < n}(G_j \cap V_j)$, the same deterministic comparison iteration as in Lemma 5.2 applies. Explicitly, $H_{\mathrm{ls}}^0 = H_{i-1}$ at a_i for $i = 1$ by identity. Equality propagates through epoch i under common maps. Then G_{i-1} equates H_{i-1} with H_i at a_{i+1}, proving the assertion inductively. The visit proxy uses the preceding boundary bits and the current bad bin. Thus its values agree as well

on every needed slab, including the earliest predecessor date of each epoch. Use V_0 in epoch zero and V_{i-1} in epoch $i \geq 1$, together with (9.7).

A physical or virtual component of diameter at least $L/100$ maps to a connected set of visited blocks of diameter at least $M/100 - 1 \geq M/200$, by the physical/coarse inequality $b(d+1)$ proved in Lemma 5.2. This contradicts the applicable V_j. The visit implication includes each intermediate factor at its original time and the final mark at a slab boundary, with the next slab initialized after all those marks.

For $t > 0$, take $n = \lceil t/(hT) \rceil$. The union bound over these failures is at most

$$CnM^4 \left(e^{-\kappa h} + 2he^{-\kappa M/200} \right).$$

Some window tests extend beyond t, which only enlarges the bounding event. With B, T, R fixed, M and h grow linearly in L; $n \leq 1 + t$. Absorb the remaining polynomial factors in a smaller positive exponential rate, as in the proof of Proposition 5.3, to obtain (9.11). At zero the syndrome is empty. There is no conditioning on earlier successful overlaps and no independence across windows. □

9.6 Pauli branches and current-syndrome recovery

Proof of Theorem 9.2. By local finiteness, all events on the finite torus can be ordered up to any finite deterministic horizon. Starting in the code space, induction over them keeps each branch a Pauli translate of its original encoded input, with the reference untouched. Every stabilizer predicate is a definite scalar on that branch and extracts no logical information. A noise string multiplies the Pauli frame regardless of its classical correlations or adaptation. By the explicit input-independence hypothesis, branch probabilities do not depend on the encoded logical amplitudes. Averaging these branches therefore defines the storage channel on encoded inputs. This argument does not assume an additive Lindblad generator for the correlated noise or apply DOA's Theorem 1 outside its Lindblad class.

Use exactly the decoder D of Section 6. Lemma 6.1 preserves the filled homology class across a single face flip if every syndrome component on both sides is small. For an atomic Pauli string, apply the lemma successively to its single-face factors in their prescribed internal order. There is no intervening correction; each partial product has the original time and fault cover. Proposition 9.6 controls all these intermediate syndromes, as well as the physical ones. It would not suffice to inspect only the two physical endpoints of an atomic error: a whole logical membrane can have empty syndrome at both endpoints. The specified order at a clock/noise tie fixes the tested syndrome, and the geometric implication covers that order.

Apply (9.11) to both marginal sectors and use a union bound. With probability at least $1 - 2C_0(1+t)e^{-\beta L}$, all virtual and physical stages through t have only small components in both sectors. Their initial filled classes are zero and remain zero across every factor. After applying the two decoder fillings, the residual Pauli is therefore a stabilizer up to phase. This branch restores all six logical qubits and any reference exactly.

The success event and its probability bound are uniform over logical inputs and allowed Pauli labels. For any pure $|\Psi\rangle_{QR}$, the target overlap is linear in the trajectory mixture and nonnegative on every remaining branch. Consequently

$$\langle \Psi | [(\mathcal{R}_L \circ \mathsf{Store}_t \circ \mathcal{E}_{\mathrm{enc}}) \otimes \mathrm{id}_R](|\Psi\rangle\langle\Psi|) | \Psi \rangle \geq 1 - 2C_0(1+t)e^{-\beta L}.$$

Take $c = \beta/3$ and enlarge C, exactly as in the proof of Theorem 2.1, to obtain (9.2) and its reference version throughout $0 \le t \le e^{cL}$. The decoder reads only the physical syndrome at the requested time. Its virtual factorization and auxiliary comparisons are proof devices, not a history supplied to recovery.

The constants can be chosen without circularity: fix the footprint geometry and $A > 100$, choose b_* for the joint witness cost, enlarge B for the clock-only term, then decrease p_* for the fixed envelope size V_B. Choose L_0 only after those choices. All are independent of the admissible law and input. □

10 Additive finite-range Markovian Lindblad noise

This section gives the strongest storage statement. The correction generator is kept exact; the added noise can couple syndrome coherences and need not preserve the identity. The stochastic-Pauli theorems retain their separate conclusions, including the correlated history model, which need not be Markovian.

10.1 Noise strength and theorem

Let

$$\mathcal{A} = \sum_p (\mathcal{D}[Z_p \Pi_p^X] + \mathcal{D}[X_p \Pi_p^Z])$$

be the unit-rate correction of (2.4), without its Pauli noise terms. Fix a microscopic interaction range r, independently of L. Let $\mathcal{N}$ be a time-independent Lindblad generator whose semigroup is completely positive and trace preserving (CPTP), with a decomposition into maps supported on qubit sets Y of diameter at most r, satisfying

$$\mathcal{N} = \sum_Y \mathcal{N}_Y, \qquad \sup_i \sum_{Y \ni i} \|\mathcal{N}_Y\|_\diamond \le \epsilon. \tag{10.1}$$

The diamond norm is the completely bounded trace norm, including an arbitrary untouched reference. The summands need not individually generate channels; their sum must be a Lindblad generator. An empty-support scalar term can be included in the trace-preserving normalization.

A Hamiltonian contribution $-i[H_Y, \cdot]$ has diamond norm at most $2\|H_Y\|$, and $\mathcal{D}[J_Y]$ has norm at most $2\|J_Y\|^2$: these follow by bounding each left/right multiplication by its operator norms, also after tensoring a reference. Thus

$$h \sum_i X_i \quad \text{has strength at most } 2|h|, \qquad \gamma \sum_i \mathcal{D}[|0\rangle\langle 1|_i] \quad \text{has strength at most } 2\gamma.$$

The coherent threshold applies to the amplitude $|h|$, not h^2. Both non-unital dissipation, such as amplitude damping, and bounded-range multi-qubit dissipation are allowed.

Theorem 10.1 (General additive-noise storage)**.** *For each fixed range r there exist $\epsilon_*(r), c, C > 0$ and an integer L_0 such that the following holds uniformly for every integer $L \ge L_0$ and every generator $\mathcal{N}$ satisfying* (10.1) *with $\epsilon \le \epsilon_*(r)$. Initialize the six logical qubits of the periodic four-dimensional toric code by ideal encoding, evolve by $e^{t(\mathcal{A}+\mathcal{N})}$, and apply the current-syndrome decoder of Section 6 in both sectors, followed by inverse encoding. For every pure state $|\Psi\rangle$ of all six logical*

qubits and any untouched reference, and every deterministic $0 \le t \le e^{cL}$*, the decoded target overlap is at least*

$$1 - Ce^{-cL}. \tag{10.2}$$

The same worst-case bound holds for any designated logical qubit with any fixed state of the unused qubits. The constants and decoder are independent of the input, reference dimension, noise generator within the bound, and read time. Preparation and final measurement and decoding are ideal.

The noise is *added* to the exact $\mathcal{A}$. Uniform changes of the correction rate are covered by time rescaling. Small spatial rate variations are also covered by using a common positive rate as baseline and treating the excess local correction dissipators as additive noise of the corresponding incident strength. No general assertion is made for replacing a correction jump $P\Pi$ by $P\Pi + \delta K$: the difference of the two dissipators need not itself be a Lindblad noise generator. General non-Markovian quantum baths are outside this theorem.

10.2 A local Pauli normal form and completely positive interpolation

Use Hermitian physical Pauli strings P_a, with $P_0 = I$. For a dissipative coefficient indexed by a, b, its *declared support* is $\operatorname{supp}(P_a) \cup \operatorname{supp}(P_b)$, even if a product $P_b P_a$ cancels some factors. Let v_r bound the number of qubits in a range-r set.

Lemma 10.2 (Local normal form)**.** *Every generator in* (10.1) *has a representation*

$$\begin{aligned}\mathcal{N}(X) = & -i\left[\sum_{a\neq 0} h_a P_a, X\right] \\ & + \sum_{a,b\neq 0} c_{ab}\left(P_a X P_b - \tfrac{1}{2}P_b P_a X - \tfrac{1}{2} X P_b P_a\right),\end{aligned} \tag{10.3}$$

where h_a *are real and* c *is positive semidefinite. Every nonzero coefficient has declared support in a range-*r *set. For a constant* C_r*, independent of* L*, the sum of absolute coefficients whose declared supports touch any specified qubit is at most* $C_r\epsilon$*.*

Proof. For a local summand on Y, use the normalized maximally entangled vector $|\Omega_Y\rangle$ and the orthonormal vectors $|P_a\rangle = (P_a \otimes I)|\Omega_Y\rangle$. The entries d^Y_{ab} of its Choi matrix in this basis are exactly its coefficients for $X \mapsto P_a X P_b$. Their moduli are at most $\|\mathcal{N}_Y\|_\diamond$, because the Choi trace norm is bounded by the diamond norm and bounds each matrix element. There are at most 16^{v_r} entries. Tensoring identity maps outside Y only adds identity Pauli factors. Sum these matrices in the global Pauli basis.

The principal nonidentity/nonidentity block $c = (d_{ab})_{a,b\neq 0}$ is positive semidefinite. Indeed, compress the Choi matrix of $e^{s\mathcal{N}}$ orthogonally to $|I\rangle$. It is positive, vanishes at $s = 0$, and its derivative there is exactly c; division by $s > 0$ and the limit prove positivity. Hermiticity preservation gives $d = d^\dagger$. Subtract the dissipative sum in (10.3). All its nonidentity/nonidentity entries cancel, leaving $GX + XG^\dagger$, with the real identity/identity entry shared between the two legs. Both the original generator and each dissipative bracket have zero trace on every input. Hence $G + G^\dagger = 0$, so the remainder is a commutator with a Hermitian operator.

Its coefficients arise only from identity-leg entries and products $P_b P_a$ belonging to the same local set Y. Drop a scalar Hamiltonian. The absolute coefficient bound, including those products and

cancellations, is therefore bounded by a fixed multiple of 16^{v_r} times the incident sum in (10.1); for example $C_r = 5 \cdot 16^{v_r}$ suffices. Products do not enlarge the declared local sets. This proves the support and strength assertions. □

Lemma 10.3 (Label deletion and interpolation)**.** *Choose $x_a \in [0,1]$, $a \neq 0$, and in* (10.3) *replace h_a by $x_a h_a$ and c_{ab} by $x_a c_{ab} x_b$. The resulting $\mathcal{N}(x)$ generates CPTP evolution. Deleting every label meeting a qubit region W makes that evolution act as identity on W. For regions W_i separated pairwise by more than r, permanently delete any family of labels and scale each remaining label meeting W_i by $z_i \in [0,1]$, leaving the others at one. Then*

$$\partial_i \partial_j \mathcal{N}(z) = 0 \quad (i \neq j),$$

and $\partial_i \mathcal{N}(z)$ is a sum of Pauli sandwiches $X \mapsto uPXQ^\dagger$, with

$$\sum |u| \leq 4 C_r \epsilon |W_i| \tag{10.4}$$

uniformly in z. Every such map acts trivially on all permanently excluded regions.

Proof. The new coefficient matrix is obtained from c by diagonal congruence, and is positive semidefinite; the Hamiltonian remains Hermitian. Thus the displayed normal form gives a Lindblad generator. If all labels meeting W are zero, every surviving operator has identity on W, including products in the anticommutators. The propagator consequently has Kraus operators acting as identity there.

Tag a Hamiltonian monomial by a, and every one of the three summands of a dissipative monomial by both a, b. A substitution $x_a = 0$ deletes precisely the exponential-series words using that label, including a tagged anticommutator product proportional to identity. Since a declared support has diameter at most r, it cannot meet two separated W_i. Every surviving generator monomial therefore depends on at most one z_i, proving the mixed-derivative identity. Its possible powers are z_i and z_i^2, whose derivatives on the unit interval have modulus at most two. Expand the commutator and dissipative bracket into Pauli sandwiches, putting all scalar phases in u. Their coefficient sums are bounded by (10.4) using Lemma 10.2 and the incident sum over W_i. Permanent deletions remain in force after differentiation. □

Lemma 10.4 (Product approximation)**.** *Partition each slab into slices of lengths $d \leq \delta$, with the actual remaining duration in a final slice. In each slice apply $e^{d\mathcal{N}}$ and every local correction channel $e^{d\mathcal{D}[P\Pi]}$ in a fixed order. For fixed L and finite t, this product converges in diamond norm to $e^{t(\mathcal{A}+\mathcal{N})}$ as $\delta \downarrow 0$. A correction channel has the exact two-Kraus representation*

$$K_0 = I - \Pi + e^{-d/2}\Pi, \qquad K_1 = \sqrt{1 - e^{-d}}\, P\Pi. \tag{10.5}$$

Proof. The successful Pauli reverses both tested check signs, so $\Pi P \Pi = 0$. On the four operator blocks associated with Π and $I - \Pi$, the active diagonal block loses a factor e^{-d}, the off-diagonal blocks gain $e^{-d/2}$, and the lost diagonal weight is transferred by P. The Kraus operators in (10.5) give exactly these four solutions of the one-term master equation. Their squared products sum to identity.

For a fixed finite system, a slice product and the full exponential have the same linear term in d. Their exponential series remainders differ by at most $O_L(d^2)$ for $d \leq 1$. All complete slice maps are CPTP and have diamond norm one. Telescoping their products therefore gives total error $O_L(t\delta)$.

Constants in this limiting argument may depend on L. The storage estimates below are uniform in the mesh once it is sufficiently fine for the fixed block side. No thermodynamic limit is interchanged with this product limit. □

10.3 Intrinsic slab violations and protected regions

Use the block geometry of Section 3, with $B \geq r + 20$, enlarged for the fixed microscopic metric, $M = \lfloor L/B \rfloor$, $b = L/M$, and $T = 1000B$. The offsets Δ, A_s are those of (3.5). For a full slab impose the following two implications in both sectors: endpoint occupation of Q_z requires initial occupation in every $z + A_s$, and any visit to Q_z requires initial occupation somewhere in U_z. Visits include every intermediate single-face factor of a noise Pauli. In a short terminal slab ending at readout impose only the visit implication.

Let W_z contain the physical supports of all checks whose syndrome edges have a vertex in either sector's U_z, with a fixed microscopic collar for the two native lattice conventions. Then

$$|W_z| \leq V_B = O_{R,r}(B^4). \tag{10.6}$$

There is an integer D_{pack}, independent of B, L, t, such that from any m specified bins in one slab one can retain at least m/D_{pack} whose W_z are separated by more than r and whose internal correction tapes are disjoint. Because $B \geq r + 20$, each collar occupies only a fixed number of coarse blocks and each retained region excludes a fixed coarse neighbourhood, so D_{pack} is independent of B. This is a greedy bounded-degree packing on the coarse lattice.

Lemma 10.5 (Uniform internal service test)**.** *There is a test of the internal correction opportunities in a protected box (no data-noise label meets W_z, so interior edges change only by internal correction updates and by boundary-face corrections toggling collar edges), for both sectors together, of failure probability at most q_B, with $q_B \to 0$ as $B \to \infty$, such that passing implies the two intrinsic implications for every allowed initial check assignment and every boundary history. For product slices the assertion holds for sufficiently small mesh δ, uniformly over that boundary history. It is used only when the internal tested check values are definite.*

Proof. Run a clean cut on full faces in U_z with the restricted initial edge assignment, coupled to the internal opportunities. Exclude a chronological chain of $m_f = 1000T$ intersecting opportunities ending within two cells of Q_z. The bound in Lemma 3.5 is $CB^4(96eT/m_f)^{m_f}$. A discrepancy from the boundary must travel through such an internal suffix: initial interior edges agree, and the boundary is at distance at least $RB - 2 > m_f + 5$. The exclusion of other changes on inner edges concerns data-noise changes, while a boundary-face correction may toggle an interior collar edge and is included as boundary input in the discrepancy trace. Boundary evolution can otherwise be arbitrary.

For a failed coordinate predecessor condition, confinement keeps the clean cut beyond the corresponding face of the target block. For the fifth condition, its maximum initial height is at most $H = \lfloor b(\sum_i z_i + 4) \rfloor$. Parity is even in the interior northeast cones. Lemma 3.2 then clears each target vertex by its two-pass schedule, with failure bound

$$(2B + 1)^4 (4 \cdot 2^{12})^{8B+2} e^{-T/2}.$$

If U_z was initially empty, the clean cut stays empty. These arguments give both implications for every initial assignment. In a terminal interval of length less than T, only the fast-chain exclusion and visit implication are used, and its probability is bounded by the full-T estimate.

For a product slice use one independent Bernoulli opportunity per internal face with parameter $1 - e^{-d}$. Each Bernoulli opportunity is the indicator of at least one mark of that face's independent rate-one Poisson process in the slice, and excluding slices with two or more marks across the entire tested internal-face family makes every fixed serialization agree with their chronological order. If that family, including both sectors, has $n_B = O_R(B^4)$ faces, the second factorial moment bounds this collision probability by $n_B^2 T\delta/2$. Choose the mesh so this is at most e^{-B}. Twice the sum of the preceding two tails, plus e^{-B}, can be used as q_B; it tends to zero. Separate protected boxes have independent unconditioned tapes. The next lemma justifies the use of definite internal predicates despite a coherent exterior. □

10.4 Exact linear histories and a one-leg filter

For a product approximant expand each $e^{d\mathcal{N}}$ in its absolutely convergent series and each normal-form coefficient into tagged Pauli sandwiches. Use a common two-outcome correction record on ket and bra. Each leg of a fine term is a Pauli word with scalar no-jump factors on definite syndromes. Serialize all its Pauli factors. For arbitrary input operators, resolve each slab boundary by independent left and right global syndrome projectors, retaining the off-diagonal blocks. These are linear resolutions, not physical measurements. For encoded input the ket and bra histories are already syndrome definite whenever the corresponding amplitude is nonzero. Absolute convergence at fixed finite L, t, δ justifies the regroupings below.

For a bin $u = (z, k)$ set $J_u = 1$ if a tagged data label on either leg meets W_z, or, when there is no such label, if the *ket* violates an intrinsic implication in at least one sector. Set $\mathrm{good}_u = 1 - J_u$. Zero-amplitude histories contribute zero. In a protected box, remaining noise acts off all tracked checks; a correction jump has a fixed Pauli shift of their signs, even if its predicate reads other signs, and a no-jump preserves them. The intrinsic mask is therefore determined by initial tracked signs and the correction record.

Write Φ_S for the linear map with exact bad-bin set S, before readout, and define its superset sum

$$\mathcal{E}(K) = \sum_{S \supseteq K} \Phi_S. \tag{10.7}$$

Superset fault-path norms are standard in quantum fault-tolerance analyses [3, 9]; the estimate below is proved for the present correction. These maps need not be positive. Summing all Φ_S gives the physical product channel, with all independent left and right syndrome resolutions included.

Lemma 10.6 (Protected-halo filter)**.** *In one slab permanently delete all data labels touching any specified collection of protected regions. Between correction gates allow exact interpolated exterior noise channels and specified Pauli sandwiches acting off those regions, with their scalar coefficients factored out. For separated protected roots $\mathcal{F}_0$, impose an intrinsic violation on the ket at each root; on any further protected roots one may impose both passing implications. The resulting linear map has diamond norm at most $q_B^{|\mathcal{F}_0|/2}$, even when the bra is unrestricted and its inserted Paulis differ from the ket's.*

Proof. Choose Kraus records for every exterior channel and for each correction gate. Let Π_s be the nonzero joint eigenspace projectors of the union of tracked checks. For a common record ω, its ket product K_ω maps Π_s into $\Pi_{s \oplus d(\omega)}$: a correction jump gives a fixed Pauli shift, each no-jump commutes with all checks, and exterior factors act off their supports. Boundary jumps may test

unknown other checks without changing this translation identity. The intermediate tracked signs are fixed by s and the jump prefix. Thus the ket acceptance mask $m(s,\omega)$ is a diagonal zero/one function of its initial check block.

With an orthonormal record register define

$$V_m = \sum_\omega |\omega\rangle \otimes \sum_s m(s,\omega) K_\omega \Pi_s, \qquad V_R = \sum_\omega |\omega\rangle \otimes K_\omega^R.$$

The unrestricted bra circuit is a composition of channels and norm-one unitary insertions, so V_R is an isometry. For fixed ω, the final translated check blocks from different s are orthogonal. Hence

$$\|V_m\|^2 = \max_s \sup_{\substack{\psi \in \operatorname{ran} \Pi_s \\ \|\psi\|=1}} \sum_\omega m(s,\omega) \|K_\omega \psi\|^2.$$

There is no multiplicity factor for possible initial check assignments. Additional passing masks only decrease this nonnegative sum.

For fixed s, ψ, sample the one-leg record distribution. At an internal face the predicate stays definite after every preceding outcome. An independent Bernoulli opportunity of parameter $1 - e^{-d}$ gives exactly the correct probabilities: accept on an opportunity when enabled, and otherwise record no jump. The normalized vector is unchanged or has the indicated Pauli applied. Sample every other Kraus outcome from its ordinary conditional probability given that vector. Pre-sampling the independent internal tapes and using fresh randomness elsewhere therefore reproduces the one-leg distribution by the chain rule, also for states entangled with the exterior.

By Lemma 10.5, a violation in every root of $\mathcal{F}_0$ requires failure of all its independent unconditioned tape tests, whatever boundary records arise. The sum is at most $q_B^{|\mathcal{F}_0|}$, uniformly in ψ and references. The filtered map is $X \mapsto \mathrm{Tr}_{\mathrm{record}}(V_m X V_R^\dagger)$. Partial trace has diamond norm one, and this sandwich has completely bounded trace norm at most $\|V_m\|\|V_R\|$. Taking the square root proves the assertion. The one-leg probability law was used only for this squared-norm bound; no fictitious rejection record has been attached to a physical no-jump amplitude. □

10.5 Mixed differences and joint diamond cost

Proposition 10.7 (Superset bound with optional faults resummed)**.** *For every $b_* \in (0,1)$, one can choose B, then a positive ϵ_*, and then a mesh bound, all independently of L, t, so that for every admissible generator of strength at most ϵ_*, every sufficiently large L, every finite read time, and every set K of real spacetime bins in its product approximant,*

$$\|\mathcal{E}(K)\|_\diamond \le b_*^{|K|}. \tag{10.8}$$

Proof. Begin with m required bins in one slab. Retain a separated subfamily of size $\ell \ge m/D_{\mathrm{pack}}$. For every other required bin expand its indicator as $J_u = 1 - \mathrm{good}_u$, giving at most $2^{m-\ell}$ signed terms. A chosen good indicator permanently deletes all labels meeting its W_z and imposes both ket passing implications. Delete the union of such labels when regions overlap; keep all passing masks. This accounts for every unselected required restriction.

At a retained root its bad indicator has two disjoint alternatives: a tagged label touches its region, or no label touches it and its ket has an intrinsic violation. Partition the retained roots into $\mathcal{I}$ of the former kind and $\mathcal{F}_0$ of the latter. Permanently delete all labels touching the clean regions,

including $\mathcal{F}_0$, and impose their masks and the previous passing masks. No intrinsic mask is imposed at $\mathcal{I}$. Scale every remaining label meeting W_i, $i \in \mathcal{I}$, by z_i as in Lemma 10.3. The product of differences between $z_i = 1$ and $z_i = 0$ selects exactly the words having a tagged appearance at every such root: a missing variable cancels, while any positive power contributes $1^a - 0^a = 1$. Both-leg tags on all anticommutator terms are retained.

This difference can be evaluated with the masks fixed: on permanently protected regions their initial values and correction record determine them, independently of exterior Kraus records and of the remaining noise-word expansion. Equivalently fix those correction records and input syndrome projectors while taking the difference and sum them afterwards.

Represent the mixed difference by the integral of the mixed derivative over $[0,1]^{|\mathcal{I}|}$. For a noise interval the exact Duhamel identity is

$$\partial_i e^{d\mathcal{N}(z)} = \int_0^d e^{(d-s)\mathcal{N}(z)} (\partial_i \mathcal{N}(z)) e^{s\mathcal{N}(z)} \, \mathrm{d}s.$$

It follows by differentiating the exponential series or by variation of constants. Iterating distinct derivatives inserts each $\partial_i \mathcal{N}$ exactly once, because $\partial_i \partial_j \mathcal{N} = 0$ for separated roots. All chronological placements are included, even within a single noise interval. For labeled insertion times their ordering simplices partition the product of $|\mathcal{I}|$ copies of the union of noise intervals. That union has total duration at most T per variable, so the integration measure is at most $T^{|\mathcal{I}|}$. The interpolation cube has volume one; correction gates between intervals use no noise-integration time.

Expand each designated derivative into Pauli sandwiches. By (10.4) their absolute coefficient sum is at most $4C_r V_B \epsilon$. Every remaining exact propagator is CPTP and acts off all permanently protected regions. The designated sandwiches also act off them, even when an extra passing region overlaps a dirty root, because its touching labels were deleted. For fixed insertion times and labels apply Lemma 10.6. The resulting partition norm is at most

$$(4C_r T V_B \epsilon)^{|\mathcal{I}|} q_B^{|\mathcal{F}_0|/2}.$$

All optional fault words have already been resummed into CPTP propagators before norms are taken. There is no factor exponential in their spacetime volume.

Sum the partitions and earlier signs. The one-slab norm is at most

$$2^m (4C_r T V_B \epsilon + \sqrt{q_B})^{\ell}.$$

Choose B large enough that $\sqrt{q_B} \le (b_*/2)^{D_{\mathrm{pack}}}/2$, then $\epsilon_* > 0$ so that $4C_r T V_B \epsilon_* \le (b_*/2)^{D_{\mathrm{pack}}}/2$. Choose the mesh small enough for the clock bound. Since $D_{\mathrm{pack}} \ell \ge m$, this gives a bound $2^m (b_*/2)^{D_{\mathrm{pack}}\ell} \le b_*^m$.

For several slabs, the restricted maps compose in time order. Their independent left and right boundary syndrome resolutions are exactly those required in the following slab; their sum preserves all cross blocks. No record is shared between slabs except through the data operator. A slab with no requirement gives its original CPTP product. Diamond-norm submultiplicativity therefore gives (10.8). The short terminal slab uses only visit masks and the same upper bounds. The order of choices leaves a strictly positive strength threshold independent of L, t. □

10.6 Walk extraction on the finite torus

Use the product metric ρ of (4.4). A unit-step walk has successive sites at distance at most one.

Lemma 10.8 (A crossing yields a short walk through many distinct marks)**.** *There are positive coarse constants a, C_0, independent of B, L, t, such that a fine ket history from the code with a syndrome component of diameter at least $L/100$, in either sector at any physical or serialized stage, supplies $k \geq aM$ distinct real bins with $J_u = 1$ and a unit-step walk of length at most $C_0 k$ visiting all those bins.*

Proof. For the chosen sector define the all-one comparisons and visit proxies by (4.1)–(4.2), using the intrinsic marks J. A good full bin has the five endpoint implications, so physical boundary occupation forces a comparison zero by induction. Every visit forces the corresponding proxy, including virtual noise stages. The partial final slab uses only this visit implication; its artificial endpoint bit need not describe the readout state.

Take $h = \lfloor M/[100(R+6)] \rfloor$, and initialize a comparison H_j at each $a_j = jh$. Extend J by zero beyond the last real slab up to the final needed overlap. Require equality of H_j, H_{j+1} at a_{j+2}, and absence of a spatially connected proxy set of diameter at least $M/200$ at every date in that two-epoch window. If every test passes, comparison equality propagates across the seams exactly as in Lemma 5.2: at the beginning of epoch one it is identity, common maps propagate it, and the next overlap equality changes the recent comparison. Predecessor proxies agree on the same dates. A large physical or virtual component would give visited block diameter at least $M/100 - 1 \geq M/200$, a contradiction.

Thus some deterministic window test fails. For an endpoint disagreement the older bit is zero and the newer one is one. Trace disagreeing predecessors back to the newer initialization date. This gives a path of older zeros, hence proxy sites, with temporal span h and steps at most R. For a spatial test failure choose a connected spatial path reaching diameter $M/200$. In either case the endpoints of a proxy path in a depth-$2h$ window have distance

$$s \geq \min\{h, M/200\}.$$

Only real marks can serve as sinks; the artificially appended future has none.

For each proxy site, a direct mark is a one-sink certificate. Otherwise take its zero predecessor and its finite backward cone down to the all-one window bottom. Every time section lifts injectively, since $4Rh < M/2$. Copy good maps, put constant-zero maps at marks and constant-one maps at the bottom and outside. The typed graph is finite, acyclic, and has nonempty predecessor sets A_s at every nonconstant node. Its root value is the maximal trajectory value. SST Theorem 22 and Definitions 17 and 21, with Eq. (2.16), give a connected five-charge contour with distinct bad sinks, avoiding the constant-one cutoff [15]. The sink images remain distinct physical bins by the time-section injectivity. With n edges per charge and m sinks, the directly checked forms (4.6) give $n \leq (1 + 8R)(m - 1)$, as in (4.7) and SST Lemmas 38 and 40. Thus no stochastic-independence assumption enters this step.

The contour and the link to its proxy root have unit-step connecting length $O_R(n+1)$, and all sinks lie in the closed ball of radius $r_i = 10(R+6)(n+1)$ about that root. This includes the bounded predecessor link when $n = 0$. Choose one ball per proxy-path site, and retain disjoint balls greedily in decreasing radius. Their triple enlargements cover the original path. A simple chain of retained centers connects its endpoint covers, with successive distances at most $3r_i + 3r_{i+1} + R$. As in Proposition 4.2, $s \leq 7 \sum_i r_i$.

Each ball supplies at least $(n+1)/(2+8R)$ distinct sinks. Different retained balls have disjoint physical sink sets, even when balls wrap. Writing k for their total number, it follows that $\sum_i r_i \leq C_R k$,

and hence $k \geq aM$ for a positive fixed a at sufficiently large M. Join successive centers by shortest unit-step paths and attach each contour's connecting graph. Its total number of edges is $O_R(k)$. A spanning-tree tour traverses each edge at most twice and visits every distinct sink in at most $C_0 k$ steps. The walk may pass through unmarked bins; it is used to locate its distinct marked keys. All constants refer only to the fixed coarse geometry. □

10.7 Dense subsets and the superset polynomial

Decrease a if necessary so $a \leq 1$, and let S_0 be the largest power of two not exceeding $aM/2$. For large M,

$$aM/4 \leq S_0 \leq aM/2.$$

For dyadic integers $1 \leq r' \leq S_0$ put

$$g(r') = \gamma \left(\sqrt{r'} + \frac{(r')^2}{S_0} \right), \tag{10.9}$$

where the positive coarse constant γ is chosen below. Call a nonempty set W of bin keys dense if every $u \in W$ has at least $g(r')$ members of W in its ρ-ball of radius r', at every such dyadic scale.

Lemma 10.9 (Extraction at a linear top scale)**.** *The marked keys of a crossing in Lemma 10.8 contain a dense subset, for a fixed $\gamma > 0$ independent of B, L, t. Every dense set has at least γS_0 members.*

Proof. Place the k distinct keys at one chosen occurrence of each on its walk. Pad the list with empty slots to a power-of-two length n, so $k \leq n \leq C_1 k$ with fixed C_1. At the dyadic scales $r' = 1, 2, \ldots, S_0$, in increasing order, discard every currently selected key in each aligned slot interval of length r' having fewer than $g(r')$ keys. The number discarded at that scale is at most $(n/r')g(r')$, and

$$\sum_{\substack{r' \leq S_0 \\ r' \text{ dyadic}}} \frac{g(r')}{r'} \leq \gamma \left(\frac{1}{1 - 2^{-1/2}} + 2 \right).$$

Choose γ small enough that C_1 times this bound is at most $1/2$. At least $k/2$ keys survive.

For a surviving key its r'-interval retained at least $g(r')$ keys at that pass. Every subsequent larger dyadic interval either misses this interval or contains it. Deleting a containing parent would delete the surviving key as well, so none of those retained keys can later be removed. Walk times in the interval differ by less than r', so their actual key locations have distance at most r'. Padding added no keys. The final set is therefore dense. At the top scale a dense set has at least $g(S_0) \geq \gamma S_0$ members, proving the last assertion. □

Let $\mathcal{U}$ be the M^4 spatial blocks times the real slabs up to readout, and $N_w = |\mathcal{U}| \geq 1$. Define $f(S) = 1$ when $S \subseteq \mathcal{U}$ contains some dense subset, and zero otherwise. This enlarged indicator may also select correctable histories.

Lemma 10.10 (Superset polynomial and dense-set count)**.** *For*

$$a_K = \sum_{S \subseteq K} (-1)^{|K| - |S|} f(S) \tag{10.10}$$

one has $f(S) = \sum_{K \subseteq S} a_K$, $a_\varnothing = 0$, *and* $|a_K| \le 2^{|K|}$. *Every* K *with nonzero coefficient is dense. The number of dense sets of size* m *is at most*

$$\exp\left(C_2 m + C_3 \frac{m}{S_0} \log(2N_w)\right), \tag{10.11}$$

where C_2, C_3 *depend only on coarse geometry.*

Proof. Interchanging the finite subset sums proves inversion: the coefficient of $f(T)$ in $\sum_{K \subseteq S} a_K$ is $\sum_{T \subseteq K \subseteq S} (-1)^{|K|-|T|}$, which is one for $T = S$ and zero otherwise. The elementary bounds and empty coefficient follow.

If a point $u \in K$ belongs to no dense subset of K, then adding or deleting u from any subset has no effect on f. Pair those two subsets in (10.10); their contributions cancel. Thus a nonzero coefficient requires every point of K to belong to a dense subset of K. Their union, equal to K, is dense: at each point one of its dense subsets supplies all the required neighbours. This proves the support assertion.

For counting, cut the circular spatial coordinates at the fixed origin, and use aligned dyadic cubes in four spatial coordinates and one time coordinate, clipped at domain ends. In each occupied side-r' cube select a representative in K. Its radius-r' ball contains at least $g(r')$ keys. A fixed key lies in such balls from at most a fixed number of cubes, since a radius-r' coordinate interval meets only boundedly many aligned side-r' intervals, also across the periodic seam. Double counting shows that the number of occupied cubes at this scale is at most $Cm/g(r')$.

At top scale S_0 at most N_w cubes can be occupied. Choosing up to $Cm/g(S_0)$ roots costs at most $(2N_w)^{C'm/S_0}$, enlarging the constant for rounding; a nonempty dense set already has $m \ge \gamma S_0$. For each occupied nonunit cube specify the subset of its 32 children that are occupied. These masks specify the unit keys exactly. The total number of such decisions is $O(m)$, since

$$\sum_{\substack{r' \le S_0 \\ r' \text{ dyadic}}} \frac{1}{g(r')} \le \gamma^{-1} \sum_{j \ge 0} 2^{-j/2} < \infty.$$

Their choices therefore cost $e^{C_2 m}$. Combining with the root choices gives (10.11). This counts arbitrary dense unions, including unions of overlapping witnesses from different epochs. □

Proposition 10.11 (Summed bad map without a volume-time loss)**.** *There are fixed* $b_*, B, \epsilon_* > 0$, *chosen in that order, such that for* $0 < t \le e^M$ *and sufficiently fine mesh, the product-channel maps satisfy*

$$\left\| \sum_S f(S) \Phi_S \right\|_\diamond \le 2e^{-\gamma S_0 \log 2}. \tag{10.12}$$

The constants are uniform in the admissible generator and read time.

Proof. The subset inversion and definition (10.7) give the exact finite identity

$$\sum_S f(S) \Phi_S = \sum_K a_K \mathcal{E}(K).$$

For $0 < t \le e^M$, $N_w = M^4 \lceil t/T \rceil$ has $\log(2N_w) \le C_4 M$ for all large M, with C_4 independent of B, since $T \ge 1$. Together with $S_0 \ge aM/4$, Lemma 10.10 bounds the number of nonzero

coefficient supports of size m by $e^{C_5 m}$, with C_5 independent of B. Choose $b_* > 0$ so $2b_* e^{C_5} \le 1/2$. Proposition 10.7 then chooses B, the noise strength threshold, and the mesh bound. Only dense supports occur and $m \ge \lceil \gamma S_0 \rceil$, so the triangle inequality gives

$$\left\| \sum_K a_K \mathcal{E}(K) \right\|_\diamond \le \sum_{m \ge \lceil \gamma S_0 \rceil} (2b_* e^{C_5})^m \le 2e^{-\gamma S_0 \log 2}.$$

No unmeasured support indicator has been assigned a physical probability, and no factor proportional to the number of possible locations or observation dates was introduced. □

10.8 Linear readout and passage to the exact semigroup

Proof of Theorem 10.1. Use the fixed decoder of Section 6. Its small-component branch fills every cycle inside its embedded coordinate box. Across a single-face flip, Lemma 6.1 makes $p + D(s) + D(s + \partial p)$ a boundary whenever both syndromes have only components of diameter less than $L/100$. This holds for all merges and splits and in both sectors.

Consider a fine ket history whose exact bad set S has $f(S) = 0$. The contrapositive of Lemmas 10.8 and 10.9 excludes a large component at every stage in either sector. Serialize each inserted Pauli on the ket into its single-face factors. Every intermediate stage was included in the intrinsic visit test. Correction jumps are themselves single-face flips; no-jumps multiply the definite syndrome leg by a scalar. Thus both filled homology classes stay zero from the encoded start. After applying the decoder the ket is the original logical ket, including its reference, up to its scalar. Phases are absorbed into that scalar.

This implication is applied to a linear cross term, not to a classical mixture. At code input the left and right syndrome projectors evaluate on their respective Pauli-translated legs. In the final measurement branch of the ket's syndrome, a different bra syndrome vanishes. If the bra syndrome matches, that bra can be arbitrary, but the decoded ket is still in the original logical ray. For the pure target $|\Psi\rangle$, put $Q_\Psi = I - |\Psi\rangle\langle\Psi|$ on the logical output and reference. Every such decoded cross term has zero expectation against Q_Ψ, because Q_Ψ annihilates its ket. No bra diagonality or independence is required.

The full product channel is trace preserving, so its decoded infidelity is its expectation of Q_Ψ. The preceding exact vanishing gives the equality of linear functionals

$$1 - F_\Psi = \mathrm{Tr}\left[Q_\Psi \left(\mathcal{R}^{(6)} \otimes \mathrm{id} \right) \left(\sum_S f(S) \Phi_S \otimes \mathrm{id} \right) (\rho_{\Psi,\mathrm{enc}}) \right],$$

where $\mathcal{R}^{(6)}$ is the physical syndrome decoder returning all six logical qubits and $\rho_{\Psi,\mathrm{enc}}$ is the encoded pure input. Its absolute value is at most the diamond norm in (10.12): the input has trace norm one, the readout is CPTP, and $\|Q_\Psi\| \le 1$. Take this equality before the norm; all coherent cross terms and the unrestricted bra normalization have already been retained.

Consequently, for $0 < t \le e^M$,

$$1 - F_\Psi \le 2e^{-\beta M}, \qquad \beta = \gamma a \log 2/4.$$

For each fixed L, t, send the mesh to zero. Lemma 10.4 and continuity of the readout and expectation give the same bound for the exact $e^{t(\mathcal{A}+\mathcal{N})}$. This takes a finite-system product limit, with all preceding estimates uniform in its mesh.

Once B is fixed, $M \geq L/(2B)$ for large L. Choose $c \leq \min\{1, \beta\}/(2B)$. Then $t \leq e^{cL}$ implies $t \leq e^{M}$, and $2e^{-\beta M} \leq 2e^{-cL}$. At zero recovery is exact; enlarging C gives the stated uniform bound. The choices were: coarse walk and density constants, then b_*, then B, then $\epsilon_*(r)$, then mesh, with L_0 chosen after the fixed geometric parameters. There is no dependence on the logical input or reference dimension.

For a designated qubit and fixed unused logical state, purify that state into an additional untouched reference and apply the all-six-qubit bound. After discarding unused outputs, the projector onto the designated pure target contains the full target projector, so its overlap cannot decrease. This proves the designated-qubit and reference claims with the same syndrome-only decoder. □

11 Scope, parameter choices and limitations

11.1 Order of constants

For the independent-noise proof, the constants can be chosen in the following order. The buffer radius $R = 2000100$ is fixed. Choose a certificate exponent $A > 100$, which suffices for the path count, and then choose $q > 0$ small enough for (4.12). Choose B sufficiently large to suppress the two clock-only terms in (3.4), and finally $\eta_0 > 0$ small enough for its noise term. This fixes $T = 1000B$, the dependence divisor D_{dep}, and all comparison constants. Only after these choices is L required to be sufficiently large for the lifted windows and epoch estimates. The resulting L_0 can be astronomically large. The proof does not give an operating threshold of practical size. For Theorem 10.1, choose the coarse walk and density constants, then the target superset cost, then B, and then the incident noise strength; the product mesh is sent to zero afterwards. The resulting threshold is astronomically small and applies to coherent amplitude $|h|$, rather than h^2.

11.2 A wrapping syndrome that does not clear

The clean-start qualification cannot be replaced by a deterministic claim that every syndrome clears in a short time. To see this, choose two different parallel coordinate lines winding in direction 1, at coordinates 0 and w in direction 2, with the other coordinates fixed and $1 \leq w < L$. Their mod-two sum is the boundary of the periodic width-w ribbon of $(1, 2)$-plaquettes between them: its interior edges cancel and the longitudinal ends close by periodicity.

This is an attainable even-degree syndrome. At every occupied vertex there is one incoming marked direction and one outgoing marked direction. There is never a pair of incoming marked directions, so every correction predicate fails. The syndrome remains frozen under any noiseless attempt sequence.

An analogous example exists for the good-map comparison. Initialize a complete coarse coordinate line with zeros and all other bits with ones. At every site on the line the four coordinate sets A_i contain the zero predecessor $d = 0$, while A_5 contains the zero predecessor $d = e_1$ along the same periodic line. Every minimum in (4.1) is therefore zero, so that line persists under every good update. Arbitrary bottom zeros cannot be charged as fresh rare failed bins.

The overlap argument avoids this issue by using an older field that itself arose from all-one data in the preceding epoch. A disagreement across its overlap forces a long path in that particular random clean-start field. Estimates are conditioned only before the older window starts, never on

an arbitrary intermediate syndrome or comparison configuration.

11.3 What is and is not implied

Drechsler, Oreg and Aharonov's own physical model is independent single-qubit Pauli noise, as specified in their Definitions 2.25, 2.37 and 7.7 [8]. Theorem 2.1 answers their time-independent question in that model. Theorem 9.2 extends the conclusion to spacetime local stochastic Pauli noise on unit time cells, including bounded-range multi-qubit and adaptive faults with a law independent of the logical input. Its description is *Lindbladian correction, trajectory-defined noise*; a general correlated law need not have an additive Lindblad noise generator.

The local-stochastic condition depends on its time grid. A shifted unit grid has parameter at most $2\sqrt{p}$, and aligned k-fold refinement has parameter at most $p^{1/k}$, as proved in Lemma 9.1. These changes are not merely a fixed rescaling of constants.

Their *fully autonomous* notion asks for protection against general local noise by a time-independent, local, translationally invariant Lindbladian [8, Section 1.5.1, Table 1]. Theorem 10.1 is the strongest result: exponential storage under any sufficiently weak additive finite-range Markovian Lindblad noise generator, with coherent and non-unital noise included. It meets their fully autonomous definition for such noise, with translation invariance when the noise is homogeneous. The stochastic-Pauli theorems retain their stated structure and correlated-history scope. The thermal Davies-memory precedent is due to Alicki and the Horodeckis [2]; the outside-equilibrium question is stated explicitly in Drechsler–Oreg–Aharonov Section 1.4.1 [8].

The added noise generator must itself generate CPTP evolution. Uniform correction-rate changes are handled by rescaling time; small inhomogeneous rate changes can be treated above a common positive baseline as additive local dissipators. General perturbations of the correction's own jump operators, including $P\Pi \mapsto P\Pi + \delta K$, remain open, as does coupling to a general non-Markovian quantum bath. The former is part of the separate superoperator-stability problem in their Section 1.8 [8]. No theorem here gives a general conversion of synchronous memories to time-independent Lindbladians or addresses faulty final syndrome measurement.

The final decoder may process the entire current syndrome. No constant-time local classical implementation of that readout is asserted. Its mathematical role is to define a fixed recovery from the present configuration and to detect an invariant homology class along small-component histories.

A logical membrane added to an initial error does not change its syndrome, so syndrome geometry cannot determine that initial logical class. From an encoded start it is known to be trivial, and the history estimate prevents its subsequent change with high probability. On $\mathbb{Z}^4$, the result instead states uniform rooted syndrome tails and vanishing low-noise density. These are the exact finite- and infinite-volume conclusions of the proof.

References

[1] Charlene Sonja Ahn. *Extending Quantum Error Correction: New Continuous Measurement Protocols and Improved Fault-Tolerant Overhead*. PhD thesis, California Institute of Technology, 2004. URL `https://thesis.library.caltech.edu/1873/1/CSAthesisv4.pdf`.

[2] Robert Alicki, Michal Horodecki, Pawel Horodecki, and Ryszard Horodecki. On thermal stability of topological qubit in Kitaev's 4D model. *Open Systems & Information Dynamics*, 17 (1):1–20, 2010. doi: 10.1142/S1230161210000023. URL https://arxiv.org/abs/0811.0033.

[3] Panos Aliferis, Daniel Gottesman, and John Preskill. Quantum accuracy threshold for concatenated distance-3 codes. *Quantum Information & Computation*, 6(2):97–165, 2006. URL https://arxiv.org/abs/quant-ph/0504218.

[4] Charles H. Bennett and G. Grinstein. Role of irreversibility in stabilizing complex and nonergodic behavior in locally interacting discrete systems. *Physical Review Letters*, 55:657–660, 1985. doi: 10.1103/PhysRevLett.55.657. URL https://doi.org/10.1103/PhysRevLett.55.657.

[5] Maury Bramson and Lawrence Gray. A useful renormalization argument. In Rick Durrett and Harry Kesten, editors, *Random Walks, Brownian Motion, and Interacting Particle Systems*, volume 28 of *Progress in Probability*, pages 113–152. Birkhäuser, 1991. doi: 10.1007/978-1-4612-0459-6_6. URL https://doi.org/10.1007/978-1-4612-0459-6_6.

[6] Shashwat Chakraborty and Jong Yeon Lee. Renormalization for robust memory without synchronization. arXiv:2610.40651v1, 2026. URL https://arxiv.org/abs/2610.40651v1.

[7] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43:4452–4505, 2002. doi: 10.1063/1.1499754. URL https://arxiv.org/abs/quant-ph/0110143.

[8] Nadav Drechsler, Yuval Oreg, and Dorit Aharonov. Autonomous quantum memories from classical stochastic processes. arXiv:2610.05589v1, 2026. URL https://arxiv.org/abs/2610.05589v1.

[9] Gesa Dünnweber, Georgios Styliaris, and Rahul Trivedi. Quantum memory and autonomous computation in two dimensions. arXiv:2601.20818v3, 2026. URL https://arxiv.org/abs/2601.20818v3.

[10] Lawrence F. Gray. Toom's stability theorem in continuous time. In Maury Bramson and Rick Durrett, editors, *Perplexing Problems in Probability*, volume 44 of *Progress in Probability*, pages 331–353. Birkhäuser, 1999. doi: 10.1007/978-1-4612-2168-5_18. URL https://doi.org/10.1007/978-1-4612-2168-5_18.

[11] Aleksander Kubica and John Preskill. Cellular-automaton decoders with provable thresholds for topological codes. *Physical Review Letters*, 123:020501, 2019. doi: 10.1103/PhysRevLett.123.020501. URL https://arxiv.org/abs/1809.10145.

[12] Sanjeev Kumar and Hendrik Weimer. Self-correction phase transition in the dissipative toric code. arXiv:2602.19288v1, 2026. URL https://arxiv.org/abs/2602.19288v1.

[13] Fernando Pastawski, Lucas Clemente, and Juan Ignacio Cirac. Quantum memories based on engineered dissipation. *Physical Review A*, 83:012304, 2011. doi: 10.1103/PhysRevA.83.012304. URL https://arxiv.org/abs/1010.2901. arXiv:1010.2901.

[14] David Ponarovsky. Local quantum fault tolerance without expansion. arXiv:2610.38847v2, 2026. URL https://arxiv.org/abs/2610.38847v2.

[15] Jan M. Swart, Réka Szabó, and Cristina Toninelli. Peierls bounds from Toom contours. arXiv:2202.10999v3, 2026. URL `https://arxiv.org/abs/2202.10999v3`. Version used for Eqs. (5.15) and (5.28). Journal version: Journal of Theoretical Probability 39, article 45 (2026), doi:10.1007/s10959-026-01508-2.

[16] A. L. Toom. Stable and attractive trajectories in multicomponent systems. In R. L. Dobrushin and Ya. G. Sinai, editors, *Multicomponent Random Systems*, volume 6 of *Advances in Probability and Related Topics*, pages 549–575. Marcel Dekker, New York, 1980.