Undecidable passive quantum storage

Weak-storage families and strong-storage temperature layouts in three dimensions

Seth Douglas Nidhal Mghirbi

Version 1.0.0, November 2026

Abstract

Can the finite description of a three-dimensional passive quantum memory decide whether it keeps a qubit as the system grows? In general it cannot. We use the exact single-qubit X, Z Davies–Metropolis bath of the stated Hamiltonian, and build every storing region from the three-dimensional memory of Balasubramanian, Davydova and Lin.

Weak storage, a lifetime tending to infinity with size, is Π^0_2-complete at each fixed low temperature in computable finite-alphabet families, even with input-independent spectra (realizations differing by a nonlocal Clifford). One fixed inhomogeneous material with infinitely many computable block gains has a set Y of rational inverse temperatures for thermal-syndrome weak storage with $0' \oplus Y \equiv_T 0''$. For strong storage, a lifetime $\exp(n^\eta)$ at every large size (n counts all qubits), one computable layout with finitely many local matrices has a Σ^0_2-complete set of storing rational inverse temperatures. Programmed omissions supply the size quantifier; its windows are astronomically cold, geometrically spaced and robust to 1% rescaling of projectors. Its preparation picks a completed block by temperature and size; at strong-storage times on infinitely many sizes, a common encoder at m designated storing centers has worst-temperature Bell overlap at most $1/2 + 1/(2m) + o(1)$. A second such layout keeps one block, fixed by size, at every temperature and stores exactly at $\beta \geq a + (b-a)\Omega$, for rational $0 < a < b$ and Chaitin's Ω; its encoder samples classical labels at the actual temperature. Conversely, for every computable interaction with rational terms the storing set is effective F_σ, so an attained lowest storing inverse temperature is left-c.e.; closed storing rays have exactly these endpoints. Halting-hardness by boundedly many temperature queries forces arbitrarily long storage/failure alternations.

Claude Opus 5.5, GPT-6.1 Sol and GPT-6 Astra assisted with this work;
the authors take responsibility for every claim (see "How this work was done").

Contents

1 The question and the results **2**

2 The storage experiment **6**

3 Uniform channel bounds **9**

4 Address families and equal spectra **12**

5 One fixed gain-catalogue material **15**

6 The mixed-scale extension of the BDL construction **19**

7 A finite-matrix strong-storage temperature layout **22**

8 Why bounded-query hardness requires alternation **28**

9 What a storing endpoint can be **30**

10 One fixed block with a prescribed strong transition **32**

11 Bounded-width commuting-Pauli chains **41**

12 Discussion **43**

How this work was done **44**

A Finite computations and reproducibility **44**

B Imported results and their scope **46**

References **46**

1 The question and the results

Can the local description of a passive three-dimensional device tell us whether increasing its size preserves a qubit for longer? The answer depends on the experiment being described. A Hamiltonian, the bath's local coupling operators and the preparation class together determine the channel. We prove undecidability for specified experiments built from a known three-dimensional quantum memory.

The constructions have elementary local couplings and computable placement rules. Their computational difficulty lies in an asymptotic question about that placement. No computation is performed by the spins during storage. In Theorems B and C ideal preparation chooses a suitable completed region according to temperature and size; in Theorem G the region is fixed by size alone. These choices are part of the statements.

There are two lifetime scales. *Weak storage* asks only that the lifetime tend to infinity through all large specimens. *Strong storage* requires a lifetime at least $\exp(n^\eta)$, for some $\eta > 0$, at every sufficiently large size; n counts all physical qubits, including auxiliary and free spins. Definition 2.1 fixes the bath, and Definitions 2.2 and 2.3 fix the encoding classes and the success threshold. Theorems A–D and G use as imports the statements of BDL v1 [2] quoted in Sections 2.1 and 6; Theorem G also imports two theorems of Cimasoni and Duminil-Copin [12].

The first result concerns infinite families indexed by a program.

Theorem A (Weak storage in address families)**.** *Under the exact bare X, Z Davies–Metropolis bath, there are a finite β_0 and computable three-dimensional address families, with uniformly bounded range, degree and qubits per cell and one fixed finite interaction alphabet, with the following property. At each fixed computable $\beta > \beta_0$, weak storage is Π_2^0-complete under computable many-one reductions, both for thermal-syndrome encoding and for arbitrary CPTP encoding. Storing instances store with thermal logical encoding; the others fail for every encoder in a size-independent time at every finite β. The same holds for a variant family whose entire finite-volume spectrum, with multiplicities, is independent of the input.*

In the equal-spectrum variant the realizations are related by a global, generally nonlocal Clifford, so equal spectra are automatic. What changes is which local operators the bath couples to the energy eigenspaces.

Temperature can be the only input.

Theorem B (A fixed material and weak-storage temperatures)**.** *There is one computable inhomogeneous three-dimensional commuting CSS material, with bounded range, degree and local dimension and an infinite computable catalogue of dyadic block gains in* $(0, 2]$*, with the following property. Under the exact bare* X, Z *Davies–Metropolis bath, its set* Y *of rational inverse temperatures with thermal-syndrome weak storage satisfies*

$$0' \oplus Y \equiv_T 0''.$$

Hence no algorithm with a halting oracle decides Y*. Storing witnesses can be chosen as thermal-syndrome encodings into one completed finite array of blocks, selected according to temperature and size.*

The thermal-syndrome class allows an outer code across the disconnected blocks of that array (Definition 2.2). Whether Y alone computes $0''$ is open. At high temperature every encoder fails in a uniform time (Lemma 5.1); Theorem B itself concerns the thermal-syndrome class.

Strong storage needs a different construction.

Theorem C (A fixed layout and strong-storage temperatures)**.** *Under the exact frequency-resolved bare* X, Z *Davies–Metropolis bath of the stated Hamiltonian, with arbitrary CPTP encoding and Bell threshold* $7/8$*, there is one computable three-dimensional layout with finitely many local matrices, and there are computable integers* $u_1 \geq 10^{100}$*,* $u_{i+1} = \lceil 5u_i/2 \rceil$*, such that:*

1. *its set of rational inverse temperatures with strong storage is* Σ^0_2*-complete;*
2. *for a* Σ^0_2*-complete set* S*, encoded in the layout by programmed omissions, every rational* $\beta \in [1.95u_i, 2.05u_i]$ *stores strongly exactly when* $i \in S$*, and this persists when each existing projector coefficient is independently rescaled by a factor in* $[0.99, 1.01]$*;*
3. *positive witnesses use thermal logical encoding into one completed block chosen according to temperature and size, and Theorem D bounds a common encoder across distinct designated yes centers.*

The level Σ^0_2 is that of the quantifier "at every sufficiently large size"; the windows are astronomically cold and geometrically spaced. Theorem G below gives a different layout, with one block fixed by size and a single prescribed noncomputable transition.

The cost of a common preparation can be measured exactly.

Theorem D (Common preparation)**.** *For the unperturbed layout and bath of Theorem C, choose* m *distinct centers* $\beta_i = 2u_i$ *with* $i \in S$*. Suppose one arbitrary CPTP encoder must serve all of them, with separate CPTP recoveries. Along the sizes* $L = 2^{s+1}$ *with* s *large, at which every installed block is complete, its worst-temperature Bell overlap at time* $\exp(n^\eta)$*, for any fixed* $\eta > 0$*, is at most*

$$\frac{1}{2} + \frac{1}{2m} + o(1).$$

For every $0 < \eta < a_0/3$*, with* a_0 *the exponent fixed in Section 7.3, the optimal common worst-temperature overlap at time* $\exp(n^\eta)$ *has limit inferior exactly* $(m + 1)/(2m)$*.*

The following obstruction explains why the construction of Theorem C has repeated temperature windows.

Theorem E (Ordered temperature queries)**.** *Fix any material or layout, one storage notion, bath, encoding class and Bell threshold. Let $Y \subset \mathbb{Q}_{>0}$ be its storing inverse temperatures. If the halting problem reduces to Y by a total computable reduction with a uniformly bounded number of truth-table queries, then Y has arbitrarily long ordered lists of alternating storing and failing temperatures. For a many-one reduction, if every queried yes temperature tolerates homogeneous Hamiltonian rescaling by each factor in $(1-\epsilon, 1+\epsilon)$, $0 < \epsilon < 1$, any alternating list of its queried temperatures of length m satisfies*

$$\beta_{\max}/\beta_{\min} \geq (1+\epsilon)^{m-1}.$$

The ratio bound assumes the exact Davies–Metropolis bath, arbitrary CPTP recovery, and unrestricted, thermal-syndrome or thermal logical preparation (Definition 2.2).

Theorem E applies bounded-query computability; its methods are due to Beigel, Beigel–Gasarch–Gill–Owings, Kummer and Jockusch [7, 8, 23, 25]. We include the diagonal argument for completeness.

What can a storing endpoint be? The next theorem holds for every computable interaction with Gaussian-rational terms, material or layout, and needs no import.

Theorem F (Storing endpoints)**.** *Take an interaction whose coordinate algorithm prints the Gaussian-rational Hermitian coefficient matrices of its local terms exactly, with all terms multiplied by one common positive number c supplied by a computable name. Use the exact bare X, Z Davies–Metropolis bath, arbitrary CPTP encoding and threshold $7/8$. Let S be its set of inverse temperatures with strong storage. On every rational compact interval $I \subset (0, \infty)$, $S \cap I$ is effective F_σ: a union of a computable sequence of effectively closed sets. Consequently:*

1. *if $t \in S \cap I$ and S contains no $\beta \in I$ with $\beta < t$, then t is left-c.e.;*
2. *if $t \in S \cap I$ and S contains no $\beta \in I$ with $\beta > t$, then t is right-c.e.;*
3. *every isolated point of $S \cap I$ is computable.*

So an attained transition can be noncomputable only in a one-sided way. The next theorem shows that this is sharp, and that strong storage needs neither repeated windows nor a temperature-dependent choice of block.

Theorem G (One fixed block, a prescribed transition)**.** *Let $0 < a < b$ be rational and let $t \in (a, b)$ be the limit of a computable sequence of rationals. Under the exact bare X, Z Davies–Metropolis bath of the stated Hamiltonian, there is a computable three-dimensional layout with finitely many local matrices, and for each size L one block B_L determined by L alone, such that:*

1. *at every $\beta > t$ the thermal logical encoding of B_L at β has Bell overlap tending to one through time $\exp(n^\eta)$ at every large size, for one $\eta > 0$ independent of β;*
2. *at every $0 < \beta < t$, along the sizes $L = 2^{s+1}$, every CPTP encoder and recovery has Bell overlap at most $1/2 + o(1)$ from time $\exp(O_\beta((\log n)^8))$ on, so strong storage fails.*

If t is left-c.e., the layout can be chosen so that item 1 holds at every $\beta \geq t$; its strong-storing set is then exactly $[t, \infty)$. This applies to $t = a + (b-a)\Omega$, with Ω Chaitin's halting probability.

Object and input	Lifetime	Preparation	Decision statement
Address family; selector	Weak	Thermal logical witness; all-encoder negative	Π^0_2-complete, including fixed spectra
Fixed gain material; rational β	Weak	Thermal syndrome in a selected finite array	$0' \oplus Y \equiv_T 0''$
Fixed finite-matrix layout; rational β	Strong, whole volume	Thermal logical in a selected block; all-encoder negative	Σ^0_2-complete by programmed omissions
Fixed-block layout; real β	Strong, whole volume	Thermal logical in a block fixed by size; all-encoder negative on cuts	Stores at every $\beta > t$ and fails below, any limit-computable t; storing set $[t, \infty)$ for left-c.e. t
Computable interaction with Gaussian-rational terms; real β	Strong	Every CPTP encoder	Effective F_σ; attained endpoints one-sided c.e.
Diagonal Hamiltonians	Constant upper time	Every CPTP encoder	No quantum storage
Bounded-width Pauli chains	Constant upper time	Syndrome diagonal; every encoder in the sign-resolving subclass	No quantum storage in these classes

Table 1: Every row uses the exact bare X, Z Davies–Metropolis bath. Among the explicit positive constructions, the fixed gain material uses infinitely many coupling values, while both strong-storage layouts use finitely many local matrices.

The thermal logical encoding samples the block's classical labels from their Gibbs law at the actual β; only the location of the block is independent of temperature. The fabrication rule computes the bond pattern of planar Ising squares inside the layout from the program for t, so the result concerns a layout, not a material. Corollary 10.6 combines Theorems F and G. For interactions covered by Theorem F, closed storing rays have exactly the left-c.e. endpoints. For the same class of interactions, testing the optimal Bell overlap at one common time scale $\exp(n^\eta_{L_s})$, with rational $\eta > 0$ and a computable increasing sequence L_s, gives exactly the limit-computable endpoints when fixed rational fidelity bounds eventually separate every rational temperature on the two strict sides, as specified in Corollary 10.6(2). Theorem D does not apply here: there, distinct registers survive at distinct temperatures, whereas here one block survives and temperature switches its barrier (Remark 10.7).

Relation to prior work. Undecidability of the spectral gap was established by Cubitt, Pérez-García and Wolf [13], and in one dimension by Bausch, Cubitt, Lucia and Pérez-García [4]. Uncomputable phase diagrams [5], a transition at an uncomputable parameter [27], and undecidable closed-system thermalization [29] are nearby precedents. The observable here is the recoverable quantum channel under a specified thermal bath. Our released *Undecidable quantum storage* [15] concerns active time-periodic circuits, with noise rate as the input.

The positive physical ingredient is Balasubramanian, Davydova and Lin's three-dimensional memory [2], abbreviated BDL. We state precisely the fixed-scale results imported from its first version. The construction of Theorem C also needs a mixed-scale extension, proved in

Address input. $w_1 \longrightarrow [Q \text{ or } C] \quad w_2 \longrightarrow [Q \text{ or } C] \quad \cdots$

The selector changes which local operators realize a block's spectrum. Unbounded accepted word lengths supply unbounded memory levels.

Temperature input. $\beta \simeq 2u_i \longrightarrow$ species $i \longrightarrow$ last completed present block

Every species uses the same microscopic penalties. Long private arms select a window. Increasing omission intervals create the failing size subsequence.

Fixed block. $\beta \longrightarrow$ Ising squares $\longrightarrow$ cylinders $\longrightarrow$ barrier of the block B_L

The squares are ordered when $\beta > t$ and disordered when $\beta < t$; ordered squares pin the checks crossed by one slab logical.

Figure 1: The three decision mechanisms. The Hamiltonian is static during storage. In the first two, box selection occurs during ideal preparation; in the third, the box is fixed by size.

Section 6. Hayden [21] proves thermal stability for a slight variant of the BDL family, and its zero-temperature perturbative stability. Yuan [34] relates single-shot correction to thermal stability. The fixed-block construction uses the critical point of planar doubly periodic Ising models found by Cimasoni and Duminil-Copin [12], classical for layered models [1], together with the cleaning argument of Bravyi and Terhal [9] and the path method of Temme [31]. Theorem F applies elementary computable analysis [32].

Reading guide. Section 2 fixes the experiment. Section 3 proves the channel bounds used in all negative branches. The address and gain constructions are in Sections 4 and 5. Readers interested in strong storage can then read Sections 6 and 7. Section 8 is independent of that construction. Section 9 proves Theorem F and uses only Section 2. Section 10 proves Theorem G; from earlier sections it uses Lemmas 2.4 and 3.1 and the counting in Lemma 6.1. The physical imports are the BDL statements, two theorems on planar Ising models used in Section 10, and two Gaussian-sampler results used only in Corollary 8.1. Section 11 proves the bounded-width no-go and states its preparation restriction explicitly.

2 The storage experiment

A computable interaction assigns finitely many qubit slots to every cell of $\mathbb{Z}^3$ and finitely supported Hermitian terms to specified slots. A coordinate algorithm terminates and prints every term meeting a given finite region. Uniform constants bound the number of slots, term diameter, terms meeting a slot and sum of their norms. The specimens are $\Lambda_L = [0, L)^3 \cap \mathbb{Z}^3$, with every term whose full support lies inside. All unused slots are physical free qubits. A finite alphabet means a fixed finite list of local matrices and bounded relative supports, up to translation. Computability places no efficiency requirement on fabrication.

Definition 2.1 (Exact bare bath)**.** For finite $H = \sum_E EP_E$ and every physical $q = X_i, Z_i$, put

$$q_\omega = \sum_{E'-E=\omega} P_{E'} q P_E, \quad r_\beta(\omega) = \min\{1, e^{-\beta\omega}\}, \quad J_{q,\omega} = \sqrt{r_\beta(\omega)}\, q_\omega.$$

The channel $\mathcal{N}_{\beta,t} = e^{t\mathcal{L}_\beta}$ has generator

$$\mathcal{L}_\beta(\rho) = -i[H, \rho] + \sum_{q,\omega} \left(J_{q,\omega} \rho J_{q,\omega}^\dagger - \tfrac{1}{2}\{J_{q,\omega}^\dagger J_{q,\omega}, \rho\} \right). \tag{2.1}$$

Each Pauli has unit attempt rate; a positive frequency raises energy. This is the exact frequency-resolved Davies–Metropolis bath throughout the paper. It is evaluated on the unperturbed stated Hamiltonian, except for the expressly allowed rescalings of its existing commuting projectors.

The weak-coupling thermal generator originates with Davies [14]. Here its finite-dimensional formula is the definition. Since $q_\omega^\dagger = q_{-\omega}$ and $r_\beta(\omega) = e^{-\beta\omega} r_\beta(-\omega)$, pairing opposite frequencies shows directly that $\sigma_\beta = e^{-\beta H}/\operatorname{Tr} e^{-\beta H}$ is stationary. Also $\sum_\omega q_\omega^\dagger q_\omega = I$: in a fixed source eigenspace sum over the orthogonal destination eigenspaces. The total intensity of a given proposal is at most one. For commuting local terms each frequency component depends only on terms touching that proposal, and hence has bounded support. Bohr covariance moves the known Hamiltonian unitary to readout.

For a tensor sum of disjoint components, a proposal changes just one component energy. Summing all spectator energy projectors gives identity, including in the presence of total-energy degeneracies. Thus q_ω is the component frequency operator tensored with spectator identity, and the channel factorizes over components.

Definition 2.2 (Encoding classes)**.** Let Φ be the normalized Bell state of an input qubit and an untouched reference. An arbitrary encoding is a CPTP map $\mathcal{E}$ from that qubit to the specimen. An arbitrary recovery $\mathcal{R}_t$ maps the specimen back to a qubit. No side memory is retained between them. Bell overlap is

$$F(\mathcal{E}, \mathcal{R}_t; t) = \operatorname{Tr} \Phi(\mathrm{id} \otimes \mathcal{R}_t \mathcal{N}_{\beta,t} \mathcal{E})(\Phi).$$

For a commuting unfrustrated Pauli check list, *thermal-syndrome encoding*, denoted $\mathfrak{T}$, means

$$\mathcal{E}(\rho) = \bigoplus_s \pi_\beta(s) \mathcal{E}_s(\rho), \tag{2.2}$$

where π_β is the exact Gibbs syndrome law and $\mathcal{E}_s$ is any CPTP map into its eigenspace. For diagonal auxiliary labels the joint syndrome/auxiliary Gibbs law is used. A *thermal logical encoding* is the special case obtained from a framed factorization $H = I_2 \otimes H_K$ by $\rho \mapsto \rho \otimes e^{-\beta H_K}/Z_K$. A *syndrome-diagonal encoding* is any CPTP map whose outputs are block diagonal in the full check labels, also when the input has a reference. Its sector probabilities may depend on the logical input.

The class $\mathfrak{T}$ permits pure logical ancillary qubits within a fixed sector and an outer code across disconnected blocks. It is closed under tensor products, outer encoding and tracing out whole independent components. The latter follows by the product Gibbs law: condition on the retained syndrome and average the partial traces of the original sector maps over the discarded syndrome. Each resulting map is CPTP. Thermal logical encoding of one selected block and Gibbs initialization of spectators belongs to $\mathfrak{T}$.

Definition 2.3 (Weak and strong storage)**.** Fix $f = 7/8$ and an encoding class $\mathfrak{C}$. The lifetime $\tau_{\mathfrak{C}}(L, \beta)$ is the supremum of horizons T for which one encoder in $\mathfrak{C}$ admits recoveries with Bell overlap at least f for every $0 \le t \le T$. Set the lifetime to zero if no horizon is admissible. Preparation can depend on L, β, H, and recovery also on t. Weak storage at β means $\tau_{\mathfrak{C}}(L, \beta) \to \infty$ as $L \to \infty$. Strong storage means that some $\eta > 0$ and L_0 satisfy

$$\tau_{\mathfrak{C}}(L, \beta) \ge \exp(n_L^\eta) \qquad (L \ge L_0),$$

where n_L is the total number of qubits in the cube.

Terminal-time success with a fixed encoder suffices for the whole interval: at an earlier time the recovery may first apply the remaining known semigroup and then the terminal decoder.

For computability, a set is Π^0_2 if it has a form $\forall a \exists b\, R(e,a,b)$ with decidable R; it is Σ^0_2 if it has a form $\exists a \forall b\, R(e,a,b)$. Completeness refers to total computable many-one reductions. The sets $0'$ and $0''$ are respectively the ordinary halting set and halting relative to $0'$. The join $A \oplus B$ gives access to both oracles. We use the elementary effective-transcript normal forms and prove their needed instances below.

2.1 The fixed-scale BDL input

Here is the physical import from BDL v1 [2], in the convention of one energy unit per violated generator. Fix its $[[5,1,2]]$ CSS seed and a sufficiently large admissible odd scale ℓ_0. Here, with BDL's fixed geometric constants α, α_2, γ and the dilation threshold λ_* of their Lemma 3.3, an odd integer ℓ is admissible when $\ell \geq (2\alpha+1)\lambda_* + 1$; the dilation $\lambda = (\ell-1)/(2\alpha+1) \geq \lambda_*$ then makes ℓ the smallest odd integer above $(2\alpha+1)\lambda$, as BDL Section 3.2.1 requires. The abstract level-$2k$ code is

$$C_{2k} = (\mathcal{R}_Z^{(\ell_0)} \mathcal{R}_X^{(\ell_0)})^k (C_0).$$

The code is determined by the seed and the replacements; coordinates are not an input to it. The following numbered BDL statements are imported: Theorem 3.16 gives one logical qubit and a local bounded-density realization; Lemmas 3.3 and 3.15 supply the embedding existence. The size bound $n_k = \exp(\Theta(k))$ is stated in BDL Section 3.3, just before Theorem 3.16. Theorem 5.13 and Corollary 5.14 give, for its decoder and some finite b_0, the two-sector unstable-syndrome estimate; together with the size bound,

$$n_k \leq AB^k, \qquad \pi_b(B_k) \leq 4\, 2^{-2^{k-1}} \quad (b > b_0). \tag{2.3}$$

Their decoder frame transports both logical Paulis under each single-site Pauli outside B_k. We import this frame property from Definition 5.1, Lemma 5.2 and Corollary 5.3. Sections 3–5, and their indicated single-step proofs, are the scope of this import.

A computable choice of embedding follows from existence. Round a real embedding to unit cells, increasing range by at most a fixed constant and occupancy by a fixed covering factor. Keep coincident vertices in different slots. For a connected bounded-range graph with v vertices, subtract coordinate minima; each coordinate span is at most the range bound times $v-1$. Enumerate integer-cell placements in that finite box and slot assignments, and test all range and occupancy constraints. Existence guarantees termination. Disconnected bounded seed pieces can be placed adjacent with a fixed margin. Fix generous integer constants once. This uses Lemmas 3.3 and 3.15 as existence results, with a search box bounded by range times vertex count. It yields a computable increasing box width $b(k)$, uniformly bounded local Pauli matrices, and a fixed slot count.

Lemma 2.4 (Stationary thermal witness)**.** *For a completed level-$2k$ block with n_k data qubits and exact bare bath, a thermal logical encoder in its decoder frame has Bell overlap at least*

$$1 - 2n_k t\, \pi_\beta(B_k) \tag{2.4}$$

at every time t. The same bound holds with auxiliary qubits whose terms are diagonal functions of auxiliary Z's and check signs, private or shared between checks, if their joint Gibbs marginal gives the displayed unstable probability; every auxiliary jump then acts trivially

on the framed logical qubit. Under (2.3)*, the lifetime exceeds* $\exp(c2^k)$ *eventually for every fixed* $0 < c < (\log 2)/2$*.*

Proof. In each syndrome sector the Hamiltonian and the total no-jump rate act trivially on the logical coordinate. A data Pauli has a definite destination syndrome and a logical unitary between the frames. Its Bohr components only refine the auxiliary part of this move. An auxiliary Pauli preserves the logical coordinate. The partial trace over the logical coordinate therefore evolves autonomously; from the thermal preparation it stays equal to its Gibbs marginal by stationarity. This remains true with the Bell reference.

The expected number of data jumps whose source lies in B_k through time t is at most $2n_k t\pi_\beta(B_k)$. Indeed each data proposal has intensity at most one. Markov's inequality bounds the probability of any such jump by the same number. Conditional on their absence, every jump and every intervening evolution preserve the framed logical state exactly. Undoing the frame and tracing spectators recovers the Bell pair on those histories. Averaging proves (2.4) for all intermediate times. Finally the loss bound at $t = e^{c2^k}$ is at most $8AB^k \exp[(c-(\log 2)/2)2^k]$, which tends to zero. □

3 Uniform channel bounds

We need bounds that survive arbitrary encoding across all components. Convergence of each fixed component, without a quantitative tensor-product argument, would not suffice.

Lemma 3.1 (Finite blocks)**.** *For the exact bare bath on a b-qubit Hamiltonian of energy width* Δ_H*, every input with an arbitrary reference is within trace norm*

$$2^b \sigma_{\min}^{-1/2} e^{-\lambda t}, \qquad \lambda = e^{-2\beta\Delta_H}, \quad \sigma_{\min} \geq 2^{-b} e^{-\beta\Delta_H},$$

of Gibbs replacement. For at most n disjoint components, each of size at most b and width at most $2Jb$*, their joint channel is within* ξ *of replacement by*

$$t = e^{4\beta Jb} \left[\tfrac{3}{2} b \log 2 + \beta Jb + \log(n/\xi) \right], \qquad 0 < \xi < 1. \tag{3.1}$$

A single fixed block channel is entanglement-breaking once

$$t \geq \lambda^{-1} \left[\tfrac{9}{2} \log(2^b) + \tfrac{3}{2} \beta\Delta_H \right]. \tag{3.2}$$

Thus any finite list of bounded component types, with any number of copies and arbitrary CPTP encoding and recovery, loses a qubit in a common size-independent time.

Proof. Write $\langle A, B\rangle_\sigma = \operatorname{Tr} \sigma A^\dagger B$ and $G = \mathcal{L}_\beta^*$. Expansion of the Lindblad product rule followed by stationarity gives

$$-\Re\langle A, GA\rangle_\sigma = \tfrac{1}{2} \sum_{q,\omega} r_\beta(\omega) \, \|[q_\omega, A]\|_{2,\sigma}^2 . \tag{3.3}$$

The Hamiltonian part is a derivation and cancels. Every rate is at least $e^{-\beta\Delta_H}$. In the unweighted Hilbert–Schmidt norm, frequency orthogonality gives

$$\sum_\omega \|[q_\omega, A]\|_2^2 = \lim_{T\to\infty} \frac{1}{T} \int_0^T \left\| [e^{isH} q e^{-isH}, A] \right\|_2^2 \, ds.$$

For each s, conjugate A by e^{-isH} and expand it in the orthogonal Pauli basis. Each nonidentity string anticommutes with at least one physical X_i or Z_i, so the sum over q is at least $4 \inf_z \|A - zI\|_2^2$. Since $\sigma_{\min}/\sigma_{\max} = e^{-\beta\Delta_H}$, equation (3.3) is at least $\lambda \operatorname{Var}_\sigma(A)$. Differentiating the centered squared norm and integrating gives contraction $e^{-\lambda t}$ in weighted norm. Weighted-to-operator norm conversion costs $\sigma_{\min}^{-1/2}$ for tests of norm at most one.

For an input ρ_{SR}, expand $\rho = \sum_{ij} |i\rangle\langle j| \otimes \rho_{ij}$. A purification writes $\rho_{ij} = \operatorname{Tr}_A |v_i\rangle\langle v_j|$ and gives $\|\rho_{ij}\|_1 \le \|v_i\| \, \|v_j\|$. Hence $\sum_{ij} \|\rho_{ij}\|_1 \le 2^b$. Apply duality to each system matrix unit and sum these reference block norms. This proves the stated referenced bound. Replacing the components one at a time adds their errors by trace-norm contraction, proving (3.1).

Here is a finite entanglement-breaking argument. Put $D = 2^b$, and let J_t be the normalized channel Choi matrix. The matrix-unit estimate gives

$$\|J_t - I/D \otimes \sigma\|_2 \le D\sigma_{\min}^{-1/2} e^{-\lambda t}.$$

Filter locally by $\sqrt{D}(I \otimes \sigma^{-1/2})$ to obtain $I + Y$, with $\|Y\|_2 \le D^2 \sigma_{\min}^{-3/2} e^{-\lambda t}$. Expand Hermitian Y in tensor products of Hermitian Pauli strings. Orthogonality and Cauchy–Schwarz give a sum of absolute coefficients at most $D\|Y\|_2$. If this is at most one, $I + Y$ is a nonnegative sum of I and matrices $I \pm P \otimes Q$. Each is positive and diagonal in a product basis across the channel input/output cut, hence separable. Removing the local filter preserves separability. The lower bound on $\sigma_{\min}$ yields (3.2).

A separable Choi matrix $\sum_a p_a R_a \otimes S_a$ gives the measure-and-prepare channel $\rho \mapsto \sum_a \operatorname{Tr}(D p_a R_a^T \rho) S_a$; trace preservation makes its measurement operators sum to identity. Tensor products of these channels remain measure-and-prepare. Encoding and recovery cannot recreate entanglement with the reference, and a separable two-qubit state has Bell overlap at most $1/2$ by convexity and the product-vector overlap bound. □

Lemma 3.2 (Independent sufficient marks)**.** *Let $D \ge 1$ be an integer and let independent source events have probabilities $q_a > 0$, with each source demanded by at most D targets and each target demanding at most D sources of a given family. There are iid sufficient target marks of probability*

$$\left(1 - (1 - q_*)^{1/D}\right)^D, \qquad q_* = \inf_a q_a > 0,$$

using private padding. A mark implies every demanded source event. Independent source families and private target clocks multiply their mark probabilities.

Proof. First thin each source to common probability q_*; any leftover success is sampled conditionally. Represent a thinned source as the OR of D independent bits with bias $1 - (1 - q_*)^{1/D}$, assigning a separate bit to each target demanding it. Pad unused source slots privately. Each target requires its own assigned bits, and pad its demands to D by private bits of the same bias. Distinct targets use disjoint random variables, so the marks are iid. A demanded bit implies its source's OR. Sampling the original independent source schedule conditional on these source indicators preserves its law. This also applies to the event that a Poisson clock is silent. □

Lemma 3.3 (Broadcast bound)**.** *Fix an integer $k \ge 2$ and consider the iid product of flagged single-qubit channels that erase with probability p, measure in the computational basis with probability c, and transmit intact with probability u, where $p + c + u = 1$. If $p \ge (k-1)u$, this channel, and every CPTP postprocessing of it, has a k-receiver extension with equal channel*

marginals. For every encoder and separate recoveries its common Bell overlap is at most

$$\frac{1}{2} + \frac{1}{2k}. \tag{3.4}$$

Proof. With probability ku send the original qubit to one uniformly chosen receiver and erasures to the rest. With probability c, measure once and copy the classical bit to all receivers. With probability $p - (k-1)u$, erase for all. These probabilities are nonnegative, sum to one, and give the required marginal. Tensor over sites and apply the postprocessing separately to the outputs.

For recovered qubits $B_1, \ldots, B_k$ and reference R, let V_j insert the normalized Bell vector on RB_j and let $P_j = V_j V_j^\dagger$. Contracting Bell vectors shows $\left\| V_i^\dagger V_j \right\| = 1/2$ when $i \neq j$. The Gram operator satisfies

$$\sum_{ij} \langle v_i, V_i^\dagger V_j v_j \rangle \leq \sum_i \|v_i\|^2 + \tfrac{1}{2} \sum_{i \neq j} \|v_i\| \, \|v_j\| \leq \tfrac{k+1}{2} \sum_i \|v_i\|^2 .$$

It has the same nonzero spectrum as $\sum_j P_j$, whose norm is therefore at most $(k+1)/2$. The sum of the k Bell overlaps is bounded by that norm. Equal marginals prove (3.4). This familiar singlet-monogamy bound also appears in [24]. □

Theorem H (Diagonal Hamiltonians)**.** *Let H be computational-basis diagonal, with single-spin energy change at most J, and let each closed radius-two interaction-graph ball have at most D qubits. Under the exact bare X, Z Davies–Metropolis bath, for every finite β and integer $k \geq 2$, every CPTP encoder and recovery has Bell overlap at most $1/2 + 1/(2k)$ after a time depending only on (J, D, β, k). Thus neither thermal-syndrome nor unrestricted weak or strong storage occurs in any such uniformly bounded family, in any dimension.*

Proof. Put $\mathcal{D}_i = (\mathrm{id} + \mathrm{Ad}_{Z_i})/2$. The Z_i dissipator is $2(\mathcal{D}_i - \mathrm{id})$. A flip component is $X_j P_{j,\omega}$ with a diagonal source projector supported on its closed neighborhood S_j. Conjugation by any Z_i changes this component only by a sign. Thus each $\mathcal{D}_i$ commutes with all dissipators and can be moved to the original input. Move the Hamiltonian unitary to the final readout by Bohr covariance.

For $t \geq 1$, factor the dephasing clocks into two input layers: a first complete pinch of probability $q = 1 - e^{-2}$ at each site, and a second independent pinch of probability $1 - e^{-2(t-1)}$. Write C_0 for the first pinched set. At any j with $S_j \subset C_0$, its flip generator is, on this invariant algebra, a classical flip at rate $r_j(x) \geq r_* = e^{-\beta J}$. This is an identity on operator-valued diagonal blocks with any unmeasured complement and reference. Realize that generator by a fair-bit reset clock of rate $2r_*$ and a rate-one residual attempt that flips with probability $r_j - r_*$. Its off-diagonal rates sum to r_j, and the residual reads only S_j. At other centers retain the true quantum generator. Interleaving these independent CPTP clocks with the remaining local semigroup realizes the sum of generators, as follows by differentiating a short interval at zero.

In the first unit interval require for target i that its radius-two ball B_i lies in C_0, that all residual clocks in B_i are silent, and that its fair reset rings. Then no operation reads or transfers its original value before its forced reset. Unprepared generator supports miss it. This proves loss of that input jointly for any set of certified targets.

Lemma 3.2, applied to the pinch and silence families, gives iid sufficient erasures with probability

$$p = (1 - e^{-2r_*})[1 - (1-q)^{1/D}]^D [1 - (1 - e^{-1})^{1/D}]^D > 0.$$

The second independent dephasing layer on unmarked targets leaves intact probability $u = (1-p)e^{-2(t-1)}$. The physical channel is a postprocessing of the flagged product erasure/measurement/intact channel: sample the clock schedules conditional on these input-independent flags, insert a placeholder for an erased input and the measured basis state for a measured input, and run the remaining maps. The forced unread reset makes the placeholder immaterial. This reasoning holds simultaneously at all erased sites and with the reference.

Choose

$$t \geq 1 + \tfrac{1}{2}\max\{0, \log((k-1)(1-p)/p)\}.$$

Then $p \geq (k-1)u$, so Lemma 3.3 applies. All constants depend only on the quantities in the statement. □

Lemma 3.4 (Weighted syndrome support)**.** *Let $C \subset \mathbb{F}_2^m$ be a binary linear space with probability proportional to $\prod_i w_i^{s_i}$, where $0 < w_i \leq 1$. For every coordinate set A,*

$$\Pr(s_i = 1 \text{ for all } i \in A) \leq \prod_{i \in A} w_i. \tag{3.5}$$

Proof. If the specified fiber is empty there is nothing to prove. Otherwise its strings on A^c form a coset $v + C_0$, where C_0 is the outside projection of the fiber with zeros on A. For the function $f(x) = \prod_{i \notin A} w_i^{x_i}$,

$$w_i^{x_i} = \tfrac{1+w_i}{2} + \tfrac{1-w_i}{2}(-1)^{x_i}$$

has nonnegative Fourier coefficients. Summing its expansion over $v + C_0$ kills nontrivial characters on C_0. Every surviving character has sign at most its sign on the zero coset. Thus $\sum_{v+C_0} f \leq \sum_{C_0} f$. The numerator of the requested probability is $\prod_{i \in A} w_i$ times the former sum; its denominator includes the latter sum. This proves (3.5) for dependent as well as independent checks. □

4 Address families and equal spectra

4.1 A static field indexed by finite words

Fix a finite alphabet $\mathcal{A}$ and enumerate its words in length-lexicographic order as $w_0, w_1, \ldots$. Reserve for w_j a cube of width $b(|w_j|)$ from Section 2.1. Place it at $(a_j, 0, 0)$ with

$$a_0 = 1, \qquad a_{j+1} = a_j + b(|w_j|) + 2.$$

Slots, widths and margins are independent of the input. For a selector A, install the prescribed level-2$|w_j|$ BDL block when $A(w_j) = 1$; otherwise leave those slots free. Terms never cross box boundaries. Use the same fixed slots everywhere else. An open cube cuts at most one box, along its x direction. A coordinate algorithm enumerates boxes until it passes the coordinate and then evaluates the finite selector. Thus every local coefficient is computable exactly, without deciding any infinite-language question.

The selectors below have finite descriptions and terminate on every word. Their head-motion convention is part of their syntax: each nonhalting transition advances at least one head; an attempted overrun rejects. Consequently at most $2(|w| + 2) + 1$ transitions are possible for two heads, up to a fixed end convention. A one-head selector is an ordinary deterministic finite automaton.

Proposition 4.1 (Language criterion)**.** *For this address family and exact bare bath, at every fixed $\beta > b_0$ weak storage with thermal logical preparation occurs exactly when $L(A)$ is infinite. The same equivalence holds for the larger thermal-syndrome and arbitrary CPTP classes. If $L(A)$ is finite, their lifetimes are bounded uniformly in specimen size at each finite β. For one-head selectors the criterion is decidable.*

Proof. A finite alphabet has only finitely many words of bounded length. If $L(A)$ is infinite, its accepted lengths are unbounded. For each desired horizon choose a sufficiently long accepted word and use Lemma 2.4 in its completed block. That block is fully present in every later specimen. Tensor spectators with their Gibbs states and ignore them at readout. This proves divergence through all large sizes, without a rate in the overall specimen width.

If accepted lengths are bounded, every active full or clipped block has bounded size and belongs to finitely many types. Free spins are additional one-spin types. Lemma 3.1 gives one entanglement-breaking time for all their tensor products, including encoders entangled across different boxes. Their Bell overlap is then at most $1/2$.

For one head, retain states reachable from the initial state and able to reach an accepting state. The language is infinite exactly when the directed letter-transition graph on these states has a cycle. Repetition of a cycle gives arbitrarily long accepted words. Conversely any sufficiently long accepting path repeats a retained state. Finite graph search decides the question. □

4.2 Two-head verification and the arithmetic level

The assertion $L(A)$ is infinite has the form

$$\forall k \, \exists w \quad [|w| \geq k \ \wedge \ A(w) = 1]. \tag{4.1}$$

Its matrix is decidable, proving the Π_2^0 upper bound on the address class. We give the reduction explicitly, using the standard multihead history technique; see [22] for background.

Let $P(e, k, j)$ be any total computable predicate and let M_e search successively for a witness j for $k = 0, 1, \ldots$. After each success it enters a distinguished event state for exactly one transition. A one-sided single-tape universal machine can implement these searches with fixed tape and control alphabets, placing e in its initial configuration. Its events are infinite exactly when $\forall k \exists j \, P(e, k, j)$.

Encode the first t transitions as rows of a word, separated by a delimiter and followed by the right endmarker. The initial row has a fixed width b_e containing the input and a spare blank. Row j has exactly $b_e + j$ cells. Each row has one marked tape letter carrying the head and state. A transition appends exactly one blank cell at the right, interpreted with the appropriate new letter if written; the spare blank ensures the head never outruns the displayed tape. Every new cell is determined by the three neighboring old letters, with fixed left-boundary and virtual-blank conventions.

A deterministic two-head one-way automaton checks the initial row with its front head while its rear head waits at the first cell. Put the front head at the successor-row start. Advance both heads along successive rows, buffering a fixed number of letters. After reading the old column to the right, test the predicted new middle column. At the old delimiter finish the boundary test, demand exactly one additional successor cell with its predicted value, then demand the new delimiter. The rear head waits during these extra reads. The heads are now positioned to compare the next pair. Finite control checks unique marked heads and

remembers whether the final completed row has the event state. Accept exactly in that case. Malformed row lengths, delimiters, head markings, initial data or transition letters reject. All waiting uses at least one advancing head, so the syntactic termination convention holds.

Induction on rows proves that every accepted word is the unique true history ending at one event time. Conversely each event history is accepted. Fixed initial width and one-cell growth prevent a single event from acquiring infinitely many padded certificates. The selector is obtained effectively from e, and accepted word lengths are unbounded exactly for e in the given Π_2^0 set. Together with Proposition 4.1 and (4.1), this proves the first assertion of Theorem A. The positive inverse-temperature constant β_0 may be any fixed strict upper bound on b_0 in the chosen energy convention.

4.3 The bare-bath isospectral family

For each reserved block use two qubits per original site. In realization Q, the original CSS checks act on the first species and the second species is free. In realization C, replace every X-type support by a Z parity on the first species and put every original Z-type support as a Z parity on the second species. Keep the same coefficient for the corresponding check. A generator has the same physical-site support in both realizations, so the open-cut convention retains corresponding complete generators.

Let r_X, r_Z be the ranks of the retained binary check-support lists. The CSS syndrome images are independent: physical Z errors set the X syndrome and physical X errors set the Z syndrome. In Q, each consistent joint syndrome has dimension $2^{2n-r_X-r_Z}$, because independent commuting Pauli generators successively halve the Hilbert-space dimension and dependent signs are forced. In C, the fibers of the two binary parity maps have dimensions 2^{n-r_X} and 2^{n-r_Z}. Thus each named joint syndrome has exactly the same multiplicity, including after any cut. Summing $e^{-\beta E(s)}$ or any function of the named signs gives identical finite-volume equilibrium generating functions.

For completeness, equality is also seen directly by Clifford conjugation. An independent commuting Pauli list can be reduced to $Z_1, \dots, Z_r$ by symplectic elimination: take a nontrivial first generator to Z_1 with single-qubit Clifford gates and CNOTs, remove its Z_1 factors from the others, and repeat on the remaining commuting subspace. The two lists have the same linear relations and signs, so use the same basis of generator labels on each side. Reducing both lists to the same canonical list gives a Clifford taking every named generator of Q to its partner in C. It need not be local. This proves exact unitary equivalence as well as the spectrum count.

Install Q when the selector accepts and C otherwise, with the same free filler for every selector. At every L, the full spectrum is independent of the selector. Infinite acceptance again gives thermal logical weak storage in completed Q blocks. For finite acceptance, all Q components have bounded size, including their spares and cuts, and become entanglement-breaking in a common time by Lemma 3.1. Every C component, of any size, satisfies Theorem H with one common $k = 2$ time. Their tensor product has a two-receiver extension: broadcast the entanglement-breaking components, tensor the extensions, then recover. Lemma 3.3 gives Bell overlap at most $3/4$ for every encoder. This proves the isospectral assertion of Theorem A, with the same Π_2^0 language upper bound and reduction.

The negative family can therefore have arbitrarily large interacting components and exactly the same energy histogram as the positive family. The invariant being changed is the locality of the operators implementing those energies relative to the fixed bare bath.

5 One fixed gain-catalogue material

5.1 Colored check energies and the hot converse

Color the intersection graph of BDL check supports greedily with K colors, using a uniform finite bound on check weight and incidence. Give color $j \in \{0, \dots, K-1\}$ the coefficient $J_j = 1 + 2^{-j-2}$. A block of gain $\alpha > 0$ has Hamiltonian

$$H_\alpha = \alpha \sum_c J_{j(c)} (I - S_c)/2. \tag{5.1}$$

The catalogue includes these blocks, their literal open x-prefixes with inherited colors, and free spins. The colors are not recomputed on a cut.

For a proposal q, all flipped checks have distinct colors. Its frequency is the signed sum of their coefficients, times α. Distinct incident sign patterns give distinct sums. If their integer parts differ, their difference is a nonzero even integer, whereas the difference of all dyadic corrections has magnitude less than one. If their integer parts agree, the largest differing dyadic place exceeds the finite sum of all smaller differences. Thus each frequency source projector resolves every flipped check sign. This exact resolution is preserved by every positive block scaling and by cuts.

For the Gibbs syndrome law, apply Lemma 3.4 with $w_c = e^{-x J_{j(c)}}$, $x = \beta\alpha$. It bounds the probability of containing a specified support by $e^{-x|\tau|}$. The witness sets, counts and decoder are unchanged by the coefficients. The BDL counting argument, stated in Section 2.1 and expanded in Section 6, therefore gives (2.3) at a sufficiently large fixed x. This proves a thermal logical cold witness for the colored catalogue.

Lemma 5.1 (All-input hot bound for the colored catalogue)**.** *There are computable constants $v_h > 0$ and $T_h < \infty$, depending only on its local support and incidence bounds, such that any finite product of colored blocks with $0 < \beta\alpha \le v_h$, including cuts and free spins, has Bell overlap at most $3/4$ at time T_h for every CPTP encoding and recovery under the exact bare bath.*

Proof. Let J bound the sum of coefficients changed by a proposal at unit gain, and put $\zeta = e^{-Jv}$ for $x \le v$. Write $P_{q,\omega}$ for its source projectors. Since every acceptance is at least ζ, its dissipator has the exact all-input decomposition

$$\mathcal{L}_q = \zeta(\mathcal{B}_q - \mathrm{id}) + \mathcal{L}'_q, \quad \mathcal{B}_q = \mathrm{Ad}_q \prod_{c:\{q,S_c\}=0} \mathcal{D}_{S_c}, \quad \mathcal{D}_S = (\mathrm{id} + \mathrm{Ad}_S)/2. \tag{5.2}$$

The residual jumps are $\sqrt{r(\omega) - \zeta}\, q P_{q,\omega}$. Equality follows by expanding the anticommutators and using $\sum_\omega P_{q,\omega} = I$. Sign resolution identifies the source pinch with the displayed product. Every summand commutes as a channel with each $\mathcal{D}_{S_c}$, by Pauli sign covariance. Remove the final known Hamiltonian unitary.

Choose a physical proposal $q(c)$ anticommuting with each nontrivial retained check. At most a constant D_0 checks choose one proposal. In a late interval of length t_1, split each independent rate-ζ clock for $\mathcal{B}_q$ into D_0 equal clocks and assign one to every check choosing it. An assigned event inserts $\mathcal{D}_{S_c}$ at the original input. Its probability is $q_1 = 1 - e^{-\zeta t_1/D_0}$, independently for different checks. Condition only on these independent CPTP event schedules, leaving the residual Lindblad semigroups unobserved.

For an early interval of length 2, after the needed labels of a center have been pinched, its exact generator on that invariant algebra equals a classical-rate Pauli generator. Split it into

unconditional Pauli conjugations of rate ζ and residual attempts of rate $1-\zeta$. A residual attempt has Kraus operators

$$q\sqrt{(r_q-\zeta)/(1-\zeta)}, \qquad \sqrt{(1-r_q)/(1-\zeta)}.$$

They are local and their squared sum is identity. The unconditional generator is equivalently complete q-dephasing at rate 2ζ. At unprepared centers keep the exact quantum generator.

For target i, require all late check marks needed by every early center that could read i, silence of all its overlapping residual attempts, and both onsite base dephasings. Uniform local bounds give a constant D bounding source and target incidences in each family. Lemma 3.2 produces iid sufficient erasures with probability at least

$$p=(1-e^{-4\zeta})^2[1-(1-q_1)^{1/D}]^D[1-(1-e^{-2(1-\zeta)})^{1/D}]^D.$$

A certified target is fully replaced: every non-onsite early map reading it is absent, while its two complete onsite dephasings multiply to full replacement. Initial check pinches are Pauli twirls and commute with this replacement, also for several targets jointly. Conditional clock sampling therefore realizes the entire channel as postprocessing of iid flagged erasure, on every input with reference.

Choose an integer upper bound $J\geq 1$ and take

$$0<v_h\leq \min\{(10J)^{-1},(20D)^{-D}/(2J)\}, \qquad t_1\geq \tfrac{10}{9}D_0D\log(20D).$$

A positive rational v_h and integer t_1 can be chosen effectively. Then $\zeta\geq 1-Jv_h\geq 9/10$. Each label and silence factor in the displayed bound on p is at least $19/20$: write it as $(1-z)^D\geq 1-Dz$, and use respectively $z=e^{-\zeta t_1/(D_0D)}\leq (20D)^{-1}$ and $z=(1-e^{-2(1-\zeta)})^{1/D}\leq (2Jv_h)^{1/D}\leq (20D)^{-1}$. Also $e^{4\zeta}\geq 1+4\zeta>4$, so the base factor exceeds $9/16$. Hence $p>(9/16)(19/20)^2=3249/6400>1/2$. Set $T_h=2+t_1$ and apply Lemma 3.3 with $k=2$. Free spins have empty label lists and satisfy the same padded sufficient-event bound. □

5.2 A finite-resource reference optimum

Fix the thermal-syndrome class $\mathfrak{T}$ from Definition 2.2. For rational $v>0$ and integer $m\geq 2$, let $W_m(v)$ assert the existence of a finite list of catalogue blocks, with repetitions, rational effective strengths $0<x_j<v$, and $\mathfrak{T}$ encoding and CPTP recovery with Bell overlap strictly greater than $1-1/m$ at time m.

These are c.e. tests. The finite syndrome spaces are computed by binary elimination. Frequency masks are fixed under positive gain changes within a block. Matrix exponentials have effective series tails. Finite-dimensional CPTP maps have computable dense nets: bound Stinespring dimensions by input dimension times output dimension, grid matrix entries of isometries and normalize sufficiently close approximants. On the compact isometry spaces the Bell-overlap error has a computable continuity bound. Enumerating these nets and enclosing the channel entries eventually recognizes every strict witness. Restriction to a fixed syndrome sector is effective, and Gibbs weights vary continuously with positive strengths. This also proves continuity of any fixed finite thermal-syndrome experiment in its gains.

Put

$$c_m=\inf\{v\in\mathbb{Q}_{>0}: W_m(v)\}, \qquad c_\infty=\sup_{m\geq 2} c_m. \tag{5.3}$$

Tests are monotone in v, and success at a larger m implies success at every smaller m by simulation at readout. Each c_m is uniformly right-c.e. The hot bound gives $c_\infty>0$: choose

$m > T_h$ with $1 - 1/m > 3/4$. The colored cold blocks and Lemma 2.4 give $c_\infty < \infty$. Moreover

$$q < c_\infty \quad \iff \quad \exists m \, \exists v \in \mathbb{Q}_{>0} \quad [v > q \ \wedge \ \neg W_m(v)]. \tag{5.4}$$

The ordinary halting oracle decides each negated c.e. test. Thus c_∞ is left-c.e. in $0'$.

5.3 An oracle-free release rule

Let $K' = 0''$ in a fixed effective enumeration of oracle programs and define

$$a = 2 - \sum_{e \in K'} 4^{-(e+1)} \in [5/3, 2]. \tag{5.5}$$

The real a is right-c.e. in $0'$. Its separated base-four digits show that it is not $0'$-computable: after earlier digits have been removed, the remaining tail following digit e is at most one third of $4^{-(e+1)}$. A precision separating those two cases recovers membership in K'. Halting is undecidable relative to its own oracle, by the diagonal program that halts exactly when a purported predictor says it will not. In particular a is not left-c.e. in $0'$, since two one-sided approximations would compute it.

For rational q, the assertion $q > a$ is c.e. in $0'$, and can be written

$$q > a \quad \iff \quad \exists u \, \forall v \, R(q, u, v) \tag{5.6}$$

with decidable R. Indeed guess a finite accepting oracle transcript, including finite halting histories for positive answers. Check these and the ordinary computation syntactically. For a negative oracle answer, test that the queried program has not halted within v steps. A transcript passing all v has exactly the true answers; a true accepting transcript passes them all. This proves the normal form needed here without an oracle in fabrication.

Enumerate rationals in $[0, 2]$. At each grid level $k \geq 1$ and positive grid point $r = j/2^k \leq 2$, start a search for a refutation of $R(q, u, v)$ for every $u \leq k$ and each $q \leq r$ among the first k entries of that enumeration. Release (k, r) exactly once when all these finitely many searches succeed. Dovetail by ordinary steps, breaking ties lexicographically. Every $r < a$ releases at every level where it is a grid point. For each $\epsilon > 0$, only finitely many released pairs have $r \geq a + \epsilon$: choose $q_0 \in (a, a + \epsilon)$ and a true witness u_0 in (5.6). Beyond a fixed k, every such pair includes this unrefutable test.

A release installs one finite array containing k copies of each of the first k catalogue types at each positive grid gain at most r. Every finite list of types and gains in $(0, a)$ can therefore be approximated arbitrarily well, with arbitrary finite multiplicities, inside a single installed array. To see this choose a dyadic $r < a$ strictly above its largest target gain and take sufficiently fine levels. Conversely, only finitely many installed elementary blocks have gain at least $a + \epsilon$; finiteness counts their copies.

Place successive arrays and their elementary boxes along the positive x axis with constant gaps. Start a box at an x coordinate at least its width and keep its other extents nonnegative and below that width. Insert a free column after every dovetail step. A coordinate algorithm simulates until it has passed the coordinate, which always occurs. An open cube cuts at most one elementary box along x; its type is in the catalogue. All local gains are positive dyadics in $(0, 2]$, computable in finite time. There is one static inhomogeneous material, with an infinite program-printed catalogue of gains.

5.4 The sharp thermal-syndrome boundary

We prove, for the material just constructed,

$$\beta \in Y \quad \Longleftrightarrow \quad \beta a > c_\infty, \qquad \beta \in \mathbb{Q}_{>0}. \tag{5.7}$$

If $\beta a > c_\infty$, choose rational $v \in (c_\infty, \beta a)$. Every $W_m(v)$ has a strict finite witness. Approximate all its target gains $x_j/\beta < a$ inside a single installed completed array. Continuity preserves the strict overlap. After that array is fully included, it is available at every larger size. A temperature- and horizon-dependent choice of array proves thermal-syndrome weak storage. For each size one may choose a suitable completed array and fix its encoder over the proposed interval, as the definition requires.

If $\beta a < c_\infty$, choose rational v with $\beta a < v < c_{m_0}$ for some m_0. Only finitely many physical elementary blocks have $\beta\alpha \geq v$. Their complete joint channel, including their finitely many possible clipped variants, approaches fixed-state replacement in diamond norm by Lemma 3.1. Choose one time making the Bell-overlap error less than 10^{-3}, uniformly in cube size. After replacing these exceptional channels, tracing their inputs preserves $\mathfrak{T}$ by Definition 2.2; their fixed outputs can be supplied by the decoder.

Suppose unbounded 7/8 horizons existed. At some $t \geq m_0$ after that replacement time the remaining finite list, all at strengths below v, would have overlap at least $7/8 - 10^{-3}$, and hence error at most 0.126. Pauli-twirl the effective qubit channel with matching Pauli conjugations of its encoder and decoder. The averaged channel is Pauli with nonidentity probability $p \leq 0.126$. Independent copies protected by the five-qubit code have error at most

$$f(p) = 1 - (1-p)^5 - 5p(1-p)^4, \qquad f(p)/p = 10p - 20p^2 + 15p^3 - 4p^4. \tag{5.8}$$

For completeness, the code generators $XZZXI, IXZZX, XIXZZ, ZXIXZ$ commute and are independent, and their fifteen weight-one Pauli syndromes are distinct and nonzero. Syndrome recovery corrects each single error [26]. This finite fact is also checked exactly in Appendix A. The derivative of the displayed quotient is $10 - 40p + 45p^2 - 16p^3 > 0$ on $[0, 0.126]$, and its value at 0.126 is less than one. Concatenation therefore makes the error arbitrarily small. The averaged Bell bound is attained or exceeded by some fixed choice of the finitely many Pauli conjugations, so no shared random register is retained.

Tensoring sector encoders and applying an outer code preserve $\mathfrak{T}$. Simulate the remaining time at readout to obtain a witness at m_0, then approximate this finite list by rational strengths, retaining strict slack. Enough concatenation would give $W_{m_0}(v)$, a contradiction. Thus the negative thermal-syndrome lifetime is size-independent at each β below the boundary. The replacements and the hot bound hold for all encoders; the closure and amplification argument uses the thermal-syndrome class, so c_∞ is the thermal-syndrome reference optimum.

Equality $\beta a = c_\infty$ is impossible for rational β: it would make $a = c_\infty/\beta$ left-c.e. in $0'$. Put $\theta = c_\infty/a$. Equation (5.7) is the rational upper cut of the positive finite irrational θ. From Y one computes θ by bracketing and bisection. With $0'$, equation (5.4) lower-approximates $a = c_\infty/\theta$, while (5.5) upper-approximates it. The two approximations recover its digits and decide $0''$. Conversely failure $\beta a < c_\infty$ is c.e. in $0'$ by comparing an upper approximation to βa with a lower approximation to c_∞. Since equality never occurs at rational queries, $0''$ decides Y. This proves $0' \oplus Y \equiv_T 0''$, and in particular $Y \not\leq_T 0'$, completing Theorem B.

6 The mixed-scale extension of the BDL construction

The next lemma extends the fixed-scale import to a larger final scale. Its proof quotes the single-step statements of BDL v1 that it uses and checks their composition.

Lemma 6.1 (Mixed-scale lemma)**.** *Fix the BDL seed and a sufficiently large admissible odd ℓ_0. At depth $2k \geq 4$ use scale ℓ_0 in the first $2k-1$ alternating replacements and any admissible odd final scale $\ell_f \geq \ell_0$. Put $\ell_{\max} = \ell_f$. There are constants C_M, C_Δ, d, V, C_0 independent of k, ℓ_f such that the resulting code has one logical qubit, a uniformly local bounded-density three-dimensional embedding, and the BDL syndrome-transport decoder. Let*

$$M = C_M \ell_{\max}^2, \quad \Delta = C_\Delta \ell_{\max}^2, \quad Q = 1 + d^4, \quad \delta = [4V(\Delta+1)^2]^{-Q}.$$

Integer upper bounds may be used for M, Δ. A sufficient dimensionless syndrome field is

$$b_{\text{suff}} = 4 \log \frac{M}{\log(1+\delta)} + 1 \leq C_0(1 + \log \ell_{\max}). \tag{6.1}$$

For every Gibbs syndrome law proportional to $\exp(-\sum_c h_c s_c)$ with all $h_c \geq b_{\text{suff}}$, the two-sector unstable probability is at most $4\,2^{-2^{k-1}}$. Thermal logical storage under the exact bare bath, also with private commuting auxiliary registers of that marginal law, obeys Lemma 2.4. At each fixed final scale there are computable A_f and one universal $B > 1$ bounding both the decorated active size and box width by $A_f B^k$ when each check carries a fixed number of arms of length at most ℓ_f.

Proof. We use the definitions of the local cochain complexes and gluing maps in BDL Sections 3.2.1–3.2.6. A step takes the abstract parent Tanner square complex, subdivides its faces at its chosen odd scale and glues the fixed two-layer replacements along matching boundary complexes. Its formulas involve that parent and that step's subdivision only. No gluing identifies subdivisions made at different steps.

One logical qubit. The statements imported here are BDL Lemmas 3.10 and 3.13: each replacement has a cochain map inducing an isomorphism on H^1. To check that the proofs remain applicable, consider their local cohomology calculation. For $\mathcal{R}_X$, after eliminating the local differentials the relevant complex is

$$C^0_{i-1} \xrightarrow{\delta^0_{i-1}} C^1_{i-1} \xrightarrow{\Delta_2 \delta^1_{i-1}} (C^2_{i-1})^{\oplus 2} \oplus E^2, \qquad \Delta_2 z = (z, z).$$

The extra summands are in degree two. Injectivity of Δ_2 makes the kernel in degree one the original kernel and leaves the image from degree zero unchanged. For $\mathcal{R}_Z$ the degree-one map to degree two is unchanged, while the map from degree zero is $\delta^0_{i-1}\Delta'_2$ with $\Delta'_2(x,y) = x + y$ surjective; extra degree-zero summands map to zero. The degree-one quotient is again unchanged. These are exactly the calculations in BDL Eqs. (85)–(93) and (97)–(101), together with their local-complex elimination. Odd subdivision changes none of the local cohomology groups or the induced identity, diagonal and codiagonal maps. Each calculation refers to a single replacement. Compose its isomorphisms; the seed has $\dim H^1 = 1$.

Coarse graining and stability. The imported step statements are BDL Theorems 4.6 and 4.7:

$$\sigma_i = \delta_i f_i + F_{i-1}\sigma_{i-1}, \qquad |\sigma_{i-1}| \leq |\sigma_i| \ (\mathcal{R}_Z), \qquad |\sigma_{i-1}| \leq |\sigma_i|/2 \ (\mathcal{R}_X). \tag{6.2}$$

Their corrections have directed dependency radius one in parent vertices, and the resulting syndrome has radius two, as in BDL Definition 4.3 and Eqs. (124)–(126). These radii count parent incidences, not fine mesh edges.

Here are the dependencies of their proofs. In Theorem 4.6, take the parity of the syndrome on each local Y_z, correct within that same Y_z, and use the induced parent coboundary to certify realizability; each parity bit is bounded by the local syndrome weight. In Theorem 4.7, first clean Y_x, then use $H^1(C_x) = 0$ to remove its error component for the proof, then clean Y_q, and finally invert the diagonal map on Y_z. The cleaning estimate on Y_x chooses parities on the link graph of the parent vertex, pairs within the four regions of each two-layer face, and compares the two layer parities. Subdivision changes the lengths of the pairing strings, but the interface crossing contributions stay $1, 1, 2$ and the parity argument of Eqs. (177)–(186) is unchanged. The Y_q estimate uses the all-one local cocycle to choose the less occupied boundary pairs, as in Eqs. (191)–(192). That choice and its counting use the same two boundary layers at any odd scale. The final inverse map halves the doubled syndrome exactly. Thus each application of (6.2) has precisely the same adjacent-level hypotheses at its own scale.

The imported BDL Lemma 5.2 gives stability when the syndrome at level two is zero. Its proof follows a weight-one error through successive parent neighborhoods $X_{i,x}$, using the directed dependency radii just described and the cochain maps. It ends with the finite decoder at C_2. Our first two scales are fixed at ℓ_0, so that decoder and its correctable parent neighborhoods are unchanged. Pulling its trivial cohomology class back through the composed H^1 isomorphisms proves the same frame-transport statement. This checks the stability input as well as the weight contractions.

Components and witness size. The imported BDL Claim 5.5 says distinct connected components of the marked decoding history do not affect each other's coarse-grained syndromes. In its proof, a dependency is either a direct parent incidence or passes through a correction and its adjacent coboundary. Each such dependency is an edge or a two-edge connection of the marked graph, so it would connect the components. This argument uses only the radii in (6.2) and is unchanged at the final scale. BDL Lemma 5.6 then applies the weight contractions separately to each component.

BDL Lemma 5.8 states that the top witness is a realizable syndrome. Explicitly, composing (6.2) over a component gives

$$\tau_{2k} = F_{2k-1} \cdots F_2 \tau_2 + \delta_{2k} p, \quad p = \sum_{i \geq 3} F_{2k-1} \cdots F_i f_i(\tau_i),$$

where the steps run from level $2k$ down to level three. Both terms are coboundaries because every F_i is a cochain map and $\tau_2 = \sigma_2$ is a coboundary, as in the proof of BDL Lemma 5.8. This gives the removal operation used in the Gibbs estimate.

The imported BDL Lemmas 5.10 and 5.11 give respectively the minimum witness-graph size 2^{k-1} and

$$\sum_i |\tau_i| \leq 4|\tau_{2k}|. \tag{6.3}$$

To verify their scale dependence, on every other backward step the component weight halves, and on the intervening step it does not increase. A nonzero level-two witness forces at least 2^{k-1} nodes in its history, by assigning each syndrome point its incident node at the next level and summing the resulting geometric sequence. Summing the backward weights in pairs gives $2 \sum_{j \geq 0} 2^{-j} = 4$, proving (6.3). Neither calculation contains a mesh scale.

Counting. BDL Lemma 3.14 bounds vertex and edge incidences by 7 and 4 after each replacement. Its proof distinguishes bulk, edge and corner vertices; their counts are 6, at most

$4+3$, and the inherited corner degree. These counts are independent of subdivision. Hence each within-level decoding graph has degree at most a fixed d. A replacement has $O(\ell_i^2)$ local syndrome positions, and at most that many children per parent vertex; enlarging constants gives $M_i \le M$ and total decoding degree at most Δ.

The imported combinatorial statements are BDL Lemmas 5.9 and 5.12. The first counts at most $|V_2|(\Delta+1)^{2m}$ connected witness graphs of m nodes by a depth-first walk of length at most $2m$. The second assigns each marked node to a syndrome-bearing node within a fixed within-level radius, giving at least $m/(1+d^4)$ syndrome-bearing nodes. The within-level incidence bound is unchanged, and all interlevel counts are dominated by $\ell_{\max}$, so these bounds hold with $V \ge \max|V_2|$ for the two sectors; V is fixed by the first two levels.

For any contained top witness τ, removal is injective on realized syndromes and reduces weighted energy by at least $b|\tau|$. Thus its probability is at most $e^{-b|\tau|}$. This is the input of BDL Eqs. (218)–(220), now with nonuniform fields. By (6.3), sum over all local syndrome assignments of one occupied history node with weight $e^{-b|\tau_i|/4}$ to get

$$W \le (1+e^{-b/4})^M - 1.$$

For a fixed graph of m nodes there are at most 2^m choices of the occupied subset, at least m/Q of them occupied, so its total weight is at most $2^m W^{m/Q}$ when $W \le 1$. Consequently the one-sector failure probability is at most

$$V \sum_{m \ge 2^{k-1}} [2(\Delta+1)^2 W^{1/Q}]^m. \tag{6.4}$$

This is the counting inequality of BDL Eq. (234), with all scale-dependent quantities bounded by their maxima. BDL Eq. (235) records its logarithmic threshold dependence. We use the sufficient condition (6.1) directly: $(1+x)^M \le e^{Mx}$ and $e^{-b/4} < \log(1+\delta)/M$ give $W \le \delta$. The bracket in (6.4) is at most $1/(2V)$, so the geometric series is at most $2\,2^{-2^{k-1}}$. Sum the two sectors. Since $\log(1+\delta) \ge \delta/2$ for $0 < \delta \le 1$, equation (6.1) is bounded by $C_0(1+\log \ell_{\max})$.

Embedding and search. The imported BDL Lemma 3.3 constructs a perturbed parent embedding at sufficiently large dilation λ. Its proof requires, in Eq. (54),

$$\alpha^3(2\alpha^3\Gamma t)^p \lambda^{3-p} \le 1/4, \qquad p > 3.$$

With t, t' fixed, increasing λ only improves this inequality. The imported Lemma 3.15 returns occupancy

$$t = \frac{64\pi t'(2\alpha+1)^2}{\gamma(\alpha_2 - 1)},$$

independent of λ. Choose the same geometric constants at each step and the dilation $\lambda_i = (\ell_i - 1)/(2\alpha+1)$, so that ℓ_i is the smallest odd integer above $(2\alpha+1)\lambda_i$, as in Section 2.1. The geometric conditions and incidence induction therefore compose, including the larger final step. Rounding and the finite placement search of Section 2.1 give a computable integer realization. All but the final replacement multiply the cell count by a fixed constant; the last multiplies it by $O(\ell_f^2)$. A fixed number of length-ℓ_f arms per check multiplies this by a computable scale-dependent factor. A connected bounded-range realization fits in width bounded by range times vertex count. Absorb these factors into A_f and keep one B independent of ℓ_f. □

6.1 Arms at bounded density

We will attach r private chains to each final check, where r is a fixed integer. Choose the final odd scale in

$$4(M_u + 2N_u) + \ell_0 \le \ell(u) \le 4(M_u + 2N_u) + \ell_0 + 4, \quad M_u = 2^u, \quad N_u = 2^{3u}. \tag{6.5}$$

BDL's flattened final replacement lies on two triangles per parent face, with bounded coincident layers. Assign a fine check to one triangle containing it. Place its chain on a normal ray with unit spacing; private leaf bits and the r copies use a fixed number of adjacent slots. Ray length is at most $\ell(u)$.

A unit ball meeting one ray has a parent triangle within distance $\ell(u) + 1$. Every such triangle has a vertex within $O(\ell(u))$ of the ball center. Undo the final dilation and bounded displacement: that vertex was in a fixed-radius ball in the previous unit-scale embedding. Its occupancy and incidence bounds therefore bound the number of parent triangles independently of u, k. On any one triangle the relevant roots project into a fixed-radius disk, so fine occupancy bounds their number. A unit ball meets bounded length on each ray. Total qubit density is thus bounded by one constant depending on r and the fixed geometric parameters, independently of u, k. Rounding all positions changes only the constant range and occupancy bounds. Crossings use distinct qubits and introduce no interactions.

The decorated graph therefore has a computable integer embedding by the same finite search, now including all arm edges and root supports. It has uniformly bounded qubits per cell and local matrices. This density argument is needed because the number of auxiliary qubits per check grows with u.

Remark 6.2. If the exponent in Definition 2.3 may depend on the temperature, Lemma 6.1 can be avoided for the bare existence of a layout as in Theorem C. One uses the single scale $\ell(u_j)$ at every replacement of a species-j block. The embedding constants of BDL Lemmas 3.3 and 3.15 do not depend on the dilation, and the threshold (6.1) is still beaten by a fixed number of arms. The price is a growth constant, and hence an exponent, that depends on the species. Theorem C keeps one exponent bound for all designated temperatures and uses Lemma 6.1.

7 A finite-matrix strong-storage temperature layout

7.1 A temperature window from a private chain

For an integer u, attach to a check sign $S_c = (-1)^s$ a lead of $M_u = 2^u$ equality bonds ending at a sensor chain of $N_u = 2^{3u}$ spins. Lead violations cost $\log 2$ and sensor-bond violations cost $4 \log 2$. The first bond is $(I - S_c Z_b)/2$ times its cost, so no extra root copy is needed. Each sensor bit b has one private bit a. Its cost is zero for $a = 0$, $\log 2$ for $(b, a) = (0, 1)$, and $2 \log 2$ for $(1, 1)$. All terms commute. Both root values have a compatible zero-energy configuration. Use r independent arms per check and the geometry in Section 6.1.

Lemma 7.1 (Thermometer bounds)**.** *For one complete arm let Z_s be its conditional partition sum and $h = \log(Z_0/Z_1)$. Allow every existing positive projector coefficient independently to lie in* $[0.99, 1.01]$ *times its nominal value. There is a computable U such that, uniformly over*

these choices, $h \geq 0$ and, for every integer $u \geq U$,

$$1.1u \leq \beta \leq 2.9u : \quad h \geq 0.05u \log 2, \tag{7.1}$$

$$0 \leq \beta \leq 0.95u : \quad h \leq \log \frac{1 + e^{-2^{0.0405u}}}{1 - e^{-2^{0.0405u}}}, \tag{7.2}$$

$$\beta \geq 3.1u : \quad h \leq 2^{-0.069u}. \tag{7.3}$$

In particular the hot and cold fields tend to zero uniformly in their stated intervals. For r arms the fields add, and tracing the auxiliary registers gives the exact data Gibbs law proportional to $e^{-\sum_c h_c s_c}$.

Proof. Use energy units $\log 2$. Tracing a private bit gives $w_b = 1 + 2^{-\beta e_b}$, where $e_0 \in [0.99, 1.01]$ and $e_1 \in [1.98, 2.02]$. Thus $d_j = \log(w_{0,j}/w_{1,j}) \geq 0$. A ferromagnetic bond with violation weight $y \leq 1$ sends a right partition vector (v_0, v_1) to $(v_0 + yv_1, v_1 + yv_0)$. It preserves $v_0 \geq v_1$; multiplication by (w_0, w_1) does too. Induction from the free end proves $Z_0 \geq Z_1$.

A lead bond of violation weight x_ℓ has transfer eigenvalues $1 \pm x_\ell$. The ratio of magnetization to total partition sum is multiplied by $(1 - x_\ell)/(1 + x_\ell)$. Its magnitude at the sensor end is at most one. At the hot end,

$$\prod_{\ell=1}^{M_u} \frac{1 - x_\ell}{1 + x_\ell} \leq e^{-\sum_\ell x_\ell} \leq e^{-2^{u-1.01\beta}} \leq e^{-2^{0.0405u}}.$$

Converting magnetization to $\log(Z_0/Z_1)$ gives (7.2). At $\beta = 0$ the ratio is zero and the same bound holds by continuity.

For the cold bound, flip the root and every chain spin simultaneously in each partition configuration after tracing private bits. All equality-bond energies are unchanged. The log ratio of the local sensor weights changes by at most $\sum_j d_j \leq N_u 2^{-0.99\beta}$. Comparing the paired positive sums gives $h \leq N_u 2^{-0.99\beta}$, which is (7.3).

In the middle interval, write $L_{\text{even}}, L_{\text{odd}}$ for the lead partition sums at equal and opposite endpoints, and T_0, T_1 for the sensor-tail sums with initial bit fixed. The odd/even ratio is at most

$$a_u \leq \prod_\ell (1 + x_\ell) - 1 \leq 2M_u 2^{-0.99\beta} \leq 2\, 2^{-0.089u}$$

for large u. Let $W_0 = \prod_j w_{0,j}$. Then $T_0 \geq W_0$. The all-one sensor string contributes $W_0 e^{-\sum_j d_j}$ to T_1. Every string with walls has weight at most W_0 times the product of its wall weights, so all such strings contribute at most $2W_0 N_u 2^{-3.96\beta}$. Uniformly for large β,

$$d_j = \log(1 + 2^{-\beta e_{0,j}}) - \log(1 + 2^{-\beta e_{1,j}}) \geq \tfrac{1}{4}\, 2^{-1.01\beta}.$$

Indeed the second Boltzmann weight is at most half the first for large β, and the derivative of $\log(1 + x)$ on $[0, 1]$ is at least $1/2$. Thus

$$\frac{T_1}{T_0} \leq e^{-N_u 2^{-1.01\beta}/4} + 2N_u 2^{-3.96\beta}.$$

Since $Z_1/Z_0 = (a_u + T_1/T_0)/(1 + a_u T_1/T_0)$, it is at most the sum of these three bounds, namely

$$2\, 2^{-0.089u} + e^{-2^{0.071u}/4} + 2\, 2^{-1.356u}. \tag{7.4}$$

For $u \geq 2000$, divide these three bounds by $2^{-0.05u}$. They are at most $2\, 2^{-0.039u}$, $32\, 2^{-0.092u}$ and $2\, 2^{-1.306u}$, respectively; for the middle term use $e^{-x} \leq 2/x^2$. Each is less than $1/3$ and

decreases with u. For $u \geq 2000$ the elementary conditions used above also hold throughout the middle interval: $\beta \geq 2$, $\sum_\ell x_\ell \leq 1$ and $N_u 2^{-3.96\beta} \leq 1$. Taking logarithms proves (7.1). Conditional on the syndrome, all private registers are independent; multiplying their partition sums proves the exact marginal identity, irrespective of check dependencies. □

Choose fixed r and a computable starting u_1 large enough that

$$0.05r\,u\log 2 > b_{\mathrm{suff}}(\ell(u)) \quad (u \geq u_1). \tag{7.5}$$

Such a choice exists effectively by (6.1) and $\log \ell(u) = 3u\log 2 + O(1)$: both sides of (7.5) grow linearly in u, so a fixed r suffices. Lemmas 6.1, 7.1 and 2.4 now give strong block storage in the middle interval. Only data jumps count in (2.4); jumps of the fluctuating auxiliary spins act trivially on the logical frame.

7.2 Small effective fields imply referenced replacement

The converse must also exclude information stored in the arms themselves.

Lemma 7.2 (Uniform data field neighborhood)**.** *For the uniformly bounded BDL data-check incidence types there is a computable $h_* > 0$ with the following property. Fix one finite arm species j, a compact inverse-temperature interval $I \subset (0,\infty)$, and all permitted one-percent coefficient rescalings. For every complete block with $0 \leq h_c < h_*$, the exact bare-bath dissipative gap has a positive lower bound depending effectively on j, I, but not on depth. There is a computable integer $C_{I,j}$ such that its channel on every input with reference is within ξ/n of Gibbs replacement by*

$$C_{I,j}[n + \log(n/\xi) + 1] \tag{7.6}$$

for $0 < \xi < 1$ in any ambient cube of $n \geq 1$ qubits. The number h_* is independent of arm length and of species.

Proof. We give the local comparison argument. Vectorize matrices with $|A\rangle\!\rangle$ denoting the Hilbert–Schmidt vector of A. For the faithful Gibbs state σ define

$$M_q = q \otimes I - I \otimes (\sigma^{-1/2} q \sigma^{1/2})^T.$$

Then $M_q |F\sqrt{\sigma}\rangle\!\rangle = |[q,F]\sqrt{\sigma}\rangle\!\rangle$. Commutativity of the local Hamiltonian cancels all terms not touching q from its modular conjugation. Each proposal has at most a fixed number b_1 of local frequency pieces, and every rate on I is at least some $r_* > 0$. Equation (3.3) and the squared triangle inequality yield

$$-\Re\langle F, GF\rangle_\sigma \geq \frac{r_*}{2b_1} \sum_q \left\| M_q |F\sqrt{\sigma}\rangle\!\rangle \right\|^2. \tag{7.7}$$

This holds for the physical constituent-qubit bath, including coherences.

First consider only the data state $\rho_h \propto e^{-\sum_c h_c (I - S_c)/2}$. Let $P_i(h)$ be the support projector of $\sum_{q = X_i, Z_i} M_q^\dagger M_q$ for this state. It has bounded support and overlap degree z_0. At $h = 0$, P_i is the complement of the normalized Bell vector at doubled site i. For any set A of sites, the joint kernel of its M_q is exactly

$$\{|F\sqrt{\rho_h}\rangle\!\rangle : F = I_A \otimes F_{A^c}\}.$$

Invertible right multiplication proves that this kernel dimension is constant in all fields. It proves constant rank for one P_i and constant nullity for $P_i + P_j$, also on the finite patch

supporting them. The support projectors are continuous near zero, since their ranks are constant. There are finitely many abstract patch types: check weight and degree are bounded, as are the radii of these patches in the data incidence graph. Arm lengths do not enter them.

At zero field the smallest positive eigenvalue of $P_i + P_j$ for $i \neq j$ is at least one. Uniform finite-dimensional continuity gives one computable product neighborhood $|h_c| < h_*$ in which it is at least $1 - 1/(2z_0)$ for every overlapping pair. A sufficient rational neighborhood can be found by finite matrix enclosures around the zero-field matrices; constant rank excludes an unobserved new small positive eigenvalue. Squaring $P_i + P_j$ gives

$$\{P_i, P_j\} \succeq -(P_i + P_j)/(2z_0).$$

Disjoint pairs commute and have positive anticommutator. Therefore $K = \sum_i P_i$ satisfies $K^2 \succeq K/2$. The commutant calculation gives its unique kernel vector $\Omega_{\rho_h} = |\sqrt{\rho_h}\rangle\!\rangle$. Hence its gap is at least $1/2$.

Append a sum of product-state projector parents on the doubled flags, with unique product vector $|0,0\rangle$ on each flag pair and gap one. For each check use a unitary on its doubled private register, controlled by S_c on the data ket, taking that vector to $\sum_z \sqrt{p_{c,s}(z)}|z,z\rangle$. The conditional Gibbs probabilities $p_{c,s}$ are those of its full arm array. The check controls commute and private flags are disjoint; a bounded coloring gives a local circuit U. For fixed species its range is bounded by a computable multiple of ℓ_j. The conditional partition identity shows that U takes Ω_{ρ_h} times the flag product to Ω_σ. Its conjugated parent projectors K'_a have bounded overlap, unique common kernel Ω_σ, and gap at least $1/2$.

For each such projector choose a physical patch A_a containing its ket and bra support. Put $G_a = \sum_{i \in A_a, q = X_i, Z_i} M_q^\dagger M_q$. Every vector in $\ker G_a$ is $(F_{A_a^c} \otimes I)\Omega_\sigma$. It is killed by K'_a, since the latter commutes with this F and kills Ω_σ. Thus $\ker G_a \subseteq \ker K'_a$.

The comparison constant is effective and uniform over rescalings. Let H_{touch} be the sum of commuting physical terms touching A_a. In computing all these M_q, use $B = e^{-\beta H_{\text{touch}}/2}$ in place of $\sqrt{\sigma}$. Right multiplication $V_B : F \mapsto FB$ simultaneously conjugates the bare commutators to these M_q. Pauli expansion on A_a bounds the bare stack squared below by four times squared distance to its kernel. An invertible conjugation loses at most $\kappa(B)^2$, so

$$G_a \succeq g_a \Pi_{(\ker G_a)^\perp}, \qquad g_a \geq 4e^{-2\beta J|A_a|},$$

for a uniform local norm bound J with $\|H_{\text{touch}}\| \leq J|A_a|$. To justify the distance estimate, for x use $\left\|V_B L V_B^{-1} x\right\| \geq s_{\min}(V_B) \left\|L V_B^{-1} x\right\|$ and $\operatorname{dist}(V_B^{-1}x, \ker L) \geq \operatorname{dist}(x, V_B \ker L)/\|V_B\|$. Kernel inclusion gives $K'_a \preceq G_a/g_a$. Patch sizes and the number of repetitions of any physical M_q are effectively bounded by constants of order at most a computable power of ℓ_j; geometric range and occupancy give a bound by a constant times $(1 + \ell_j)^3$ for patch volumes. Summing this comparison and using (7.7) proves the positive depth-independent dissipative gap. It is uniform on I and on the allowed rescaling box. Frequency collisions cannot spoil the estimate: the number of signed local energy changes remains uniformly bounded, and the commutator inequality uses their sum.

Finally apply the matrix-unit and reference conversion in Lemma 3.1, now with this fixed gap. The energy width is $O(n)$, so $-\log \sigma_{\min} = O_I(n)$. This yields (7.6). Every constant used admits an effective finite overestimate from the incidence bounds, I and ℓ_j. □

7.3 Species, omissions and complete cuts

Put $u_{i+1} = \lceil 5u_i/2 \rceil$. Choose $u_1 \geq 10^{100}$ beyond Lemma 7.1, equation (7.5), and the bounds ensuring $rh < h_*$ throughout the off-target hot and cold intervals. At any $\beta \in I_i =$

$[1.95u_i, 2.05u_i]$, species $j > i$ is hot because $2.05u_i < 0.95u_j$, and species $j < i$ is cold because $1.95u_i > 3.1u_j$. Lemma 7.2 gives computable uniform integers C_{ij} for complete off-target species $j \neq i$ on I_i.

Fix a many-one complete Σ^0_2 set

$$S = \{i : \exists a \, \forall v \, P(i, a, v)\},$$

with total decidable P. Concretely take $S = \mathrm{FIN} = \{i : W_i \text{ is finite}\}$ for a standard c.e. enumeration, and let $P(i, a, v)$ say that at most a distinct numbers have entered W_i by stage v. This predicate is uniformly decidable by bounded simulation. To prove completeness, start with any decidable $R(e, a, v)$ in a Σ^0_2 definition. At stage s compute the least $a \leq s$ passing all tests $v \leq s$, using $s + 1$ if none passes, and enumerate a fresh number whenever this candidate strictly increases. Candidates never decrease. They eventually stabilize exactly when a true witness exists; otherwise they tend to infinity. The resulting c.e. set is finite exactly for the given Σ^0_2 inputs, and its program index is computable from e. Thus FIN is many-one complete. At stage s, for $i \leq s$ let $a_i(s)$ be the least $a \leq s$ passing all tests $v \leq s$, or $s + 1$ if none does. This sequence never decreases. It stabilizes exactly when $i \in S$: a true witness eventually becomes least after all smaller false ones are refuted; in its absence every fixed candidate is eventually refuted.

At the dth strict increase of $a_i(s)$, omit species i at all stages in $[s, (d + 1)s]$. Union overlapping omission intervals. Whether a stage is omitted is decided by finite simulation. Let $A_j B^k$ be the size and width bounds from Lemma 6.1. Choose a rational $a_0 > 0$ with $a_0 \log_2 B \leq 1/4$. At stage s, admit only the first $b(s) \leq s$ species, where $b(s)$ is the largest initial segment satisfying

$$\max_{j \leq b(s)} A_j \leq 2^{s/8}, \qquad \max_{i,j \leq b(s), i \neq j} C_{ij} \leq 2^{s/8}. \tag{7.8}$$

Integer eighth powers decide the comparisons. The sequence is computable, nondecreasing and tends to infinity.

Proposition 7.3 (Whole-volume layout estimates)**.** *With these choices, install one species-i block of depth $2\lfloor a_0 s \rfloor$ at every admitted nonomitted stage s. All blocks fit, with constant gaps, along x in $[2^s, 2^{s+1})$ for all sufficiently large s, and in the nonnegative other coordinates below 2^s. Use q_0 slots per cell everywhere. Then $n_L = q_0 L^3$, and under the exact bare bath:*

1. *if $i \in S$, every $\beta \in [1.1u_i, 2.9u_i]$ has thermal logical strong storage through $\exp(n_L^\eta)$ at every large L, for each fixed $0 < \eta < a_0/3$;*
2. *if $i \notin S$, every $\beta \in I_i$ has a sequence of complete cubes on which every CPTP encoder is within fixed arbitrarily small error of input-independent replacement by time $\exp(n_L^{o(1)})$.*

Both statements are uniform under independent one-percent rescalings of the existing projectors. The nominal layout has only finitely many local matrices, with penalties $\log 2, 2 \log 2, 4 \log 2$.

Proof. By (7.8), each block has active size and width at most $2^{s/8} B^{a_0 s} \leq 2^{3s/8}$. At most s blocks occur. Thus $s(2^{3s/8} + 2) < 2^s$ beyond a computable finite stage. Omit earlier stages if needed. The embedding and arm proof gives one uniform slot count, range and incidence bound. The finite list of root, equality-bond and private-bit projectors uses bounded physical supports; rounding yields only finitely many relative support patterns. Fabrication at a coordinate requires only a finite stage simulation. The bath acts independently on all component and filler qubits. The fabrication rule never receives the queried inverse temperature.

If $i \in S$, its omission intervals end after finitely many stages, and it is eventually admitted. For $2^s \le L < 2^{s+1}$ the previous stage has a complete species-i block of level $k \ge a_0(s-1)-1$. The thermometer middle bound and (7.5) give the mixed-scale unstable estimate uniformly over rescalings. Lemma 2.4 gives e^{c2^k} storage, whose logarithm dominates n_L^η for $\eta < a_0/3$. Its data-attempt prefactor is only exponential in k for fixed species. Every qubit of the cube is included in n_L; spectators factor and are ignored at readout.

If $i \notin S$, let s be the stage of its dth increase and put $s' = (d+1)s$. At $L = 2^{s'+1}$ every included box is complete and no next-stage box has entered. Every species-i block comes from a stage before s, so its size is at most $2^{3s/8} + O(1)$. Since $n_L = q_0 2^{3(s'+1)}$, this is $n_L^{o(1)}$ as $d \to \infty$. Lemma 3.1 replaces all such blocks to total error ξ by a time with

$$\log(1+t) \le O_{\beta,J}(2^{3s/8} + \log\log(n_L+2)) = n_L^{o(1)}.$$

For every other present species j, eventually $i, j \le b(s')$, so $C_{ij} \le 2^{s'/8}$. Its replacement time (7.6), using individual error ξ/n_L, is polynomial in n_L. Free spins also replace in $O(\log(n_L/\xi))$ time. Replace all component channels one by one. Errors add even with arbitrary joint encoders and references. The resulting channel is a constant; its recovered Bell overlap is $1/4$. Fixed small ξ defeats $7/8$ and every exponent $\eta > 0$ along these cuts. Uniform coefficient and energy bounds make every estimate uniform over the allowed rescalings. □

7.4 The upper arithmetic bound

The map $i \mapsto 2u_i$ reduces S to the strong-storage temperatures of the nominal layout by Proposition 7.3. To classify the entire rational set, its finite-volume optimum $F_L(\beta, t)$ over arbitrary CPTP encoders and recoveries is uniformly computable. The nominal Hamiltonian is $\log 2$ times a sum of commuting integral-spectrum projectors. Exact finite-dimensional algebra determines its joint projectors and integer frequency coincidences; at rational β the Metropolis rates are algebraic powers of two. The known unitary is computable as well. Compact Stinespring nets with certified finite-channel approximations compute the optimum with any prescribed error, as in Section 5.

Terminal success is sufficient by Definition 2.3. Compactness of the finite encoder and recovery spaces and continuity in time make successful terminal times closed, so the lifetime supremum causes no endpoint ambiguity. Replacing any successful exponent by a smaller reciprocal integer gives the equivalent predicate

$$\exists a \ge 1\, \exists N\, \forall L \ge N: \quad F_L(\beta, \exp(n_L^{1/a})) \ge 7/8. \tag{7.9}$$

A non-strict inequality for a computable real is a universal test on rational approximations: use approximants with error 2^{-p} and require each to be at least $7/8 - 2^{-p}$. Pair the precision and size quantifiers into one universal variable. Equation (7.9) is Σ_2^0. This proves Theorem C for the nominal computable layout. Arbitrary real rescalings need no effective description for the physical label inequalities, and are not assigned an arithmetic classification.

7.5 The necessity of choosing a block

We prove Theorem D. At the sizes $L = 2^{s+1}$, s large, every installed block is complete. There let A_i consist of all blocks of one designated yes species. At its center $\beta_i = 2u_i$, every component in A_i^c replaces its input in polynomial time by Lemma 7.2 and (7.8). Use individual trace error $1/n_L^2$ and sum over at most n_L components. The complement is $o(1)$

from replacement on every referenced input. This remains true at the much longer time $e^{n_L^{\eta}}$, by contraction. For finitely many centers it holds simultaneously.

Given a common encoder, the recovered output for temperature i can therefore be simulated from A_i alone: evolve that register, append the known Gibbs output of its complement and apply that temperature's decoder. The m disjoint registers support these simulations in parallel and give one joint state with the original Bell reference. The Gram proof of Lemma 3.3 bounds the sum of its overlaps by $(m+1)/2$. The actual outputs differ by $o(1)$, proving the upper bound on the common worst-temperature overlap. The same proof permits different times that all exceed the common replacement polynomial.

For attainment use the symmetric qubit cloner [18, 33], in the explicit form

$$\mathcal{C}_m(\rho) = \frac{2}{m+1}\Pi_{\mathrm{sym}}(\rho \otimes I^{\otimes(m-1)})\Pi_{\mathrm{sym}}.$$

This map is CP. Unitary invariance and $\operatorname{Tr}\Pi_{\mathrm{sym}} = m+1$ give $\operatorname{Tr}_{2\dots m}\Pi_{\mathrm{sym}} = (m+1)I/2$, so it is trace preserving. For input $|0\rangle$, its symmetric occupation state with k zeroes has probability $2k/[m(m+1)]$. One output has zero probability

$$\sum_{k=0}^{m} \frac{2k^2}{m^2(m+1)} = \frac{2m+1}{3m}.$$

Covariance makes the marginal depolarizing. A depolarizing channel of pure-state fidelity p has Bell overlap $(3p-1)/2$, giving $(m+1)/(2m)$. At every large cube encode clone i in a completed species-i block, with that block's thermal preparation for β_i. Perform all these preparations regardless of which bath will be used. Under bath i, its own block decodes as identity with $o(1)$ failure probability through the required time, also for a clone entangled with the others. Thus the common worst overlap tends to the cloning value from below. The all-size lower bound and the upper bound at the sizes $L = 2^{s+1}$ give exactly the claimed limit inferior.

In particular a common encoder with Bell margin $f - 1/2 > 0$ can serve at most $\lfloor 1/(2f-1) \rfloor$ designated centers with all-size strong guarantees. For the layout of Theorem C the block choice is therefore essential. Remark 10.7 explains why the fixed-block layout of Theorem G is not subject to this bound.

8 Why bounded-query hardness requires alternation

We prove Theorem E from a finite counting argument. It applies to the weak and strong experiments alike, once the material or layout, bath, preparation class and success convention are fixed. Its computability content is part of the theory of query-limited reducibility and verbose sets [7, 8, 25]; cuts of computable linear orders connect it with semirecursive sets [23].

Suppose first that $Y \subset \mathbb{Q}_{>0}$ is the trace of a union of at most h real intervals. On m sorted distinct query points a membership vector has at most h runs of ones. Describe each run by its first and last discrete positions, padding with empty runs. At most $(m+2)^{2h}$ possible vectors result. This finite superset of transcripts is effectively enumerable without knowing any interval endpoints.

Assume the halting set $K = \{e : \varphi_e(e) \downarrow\}$ reduces to Y using at most d truth-table queries per input. For adaptive total query trees of depth d, expand every possible transcript beforehand; at most $D = 2^d - 1$ distinct query nodes per input suffice. For a nonadaptive

reduction put $D = d$. Choose N with

$$2^N > (DN + 2)^{2h}. \tag{8.1}$$

We recall the simultaneous self-reference needed here and prove it. In an acceptable program enumeration let $s(z, w)$ be the computable operation that hard-codes the first argument of a two-argument program. For any total code transformation T, take a two-argument program $v(z, x)$ that executes code $T(s(z, z))$ on x. Then $u = s(v, v)$ computes the same function as $T(u)$, by substitution. To obtain N mutually referring programs, make $T(u)$ a program on parameter i that knows the complete list of codes for "execute u on j and ignore the external input," $j = 1, \ldots, N$. Harmless distinct literals ensure distinct codes. The fixed point makes their behavior exactly the requested common-list behavior. Their halting on their own code is the same as their parameterized nullary halting.

Let these codes be $e_1, \ldots, e_N$. Each program computes all the reduction's query trees and answer tables on that common list, forms the sorted union of at most DN query points, and enumerates every membership vector with at most h runs. It computes the N predicted halting answers for each such vector, including the appropriate adaptive path. By (8.1), an N-bit vector is absent from the predictions. Choose the lexicographically first absent vector. Program e_i halts exactly when its bit is one. No call to Y occurs in this construction, and every finite computation terminates. The actual Y transcript is among those enumerated, yet its reduction prediction must equal the omitted halting vector. This contradiction excludes every bounded truth-table reduction.

If a subset of a linearly ordered set has a bound on the length of alternating membership lists, it has only finitely many convex yes components: choose a yes point in each alleged component and intervening no points to obtain longer alternating lists. Each convex subset of the positive rationals is the trace of a real interval with suitable endpoint inclusions. Thus absence of arbitrary alternation would give the finite-interval case just excluded. The same proof works on any prescribed query subinterval, or with an interval extension of promised labels. The ray in Theorem B raises no contradiction: its relativized Turing lower bound is a different assertion from bounded-query halting hardness.

8.1 Calibration addendum

For a total many-one reduction $r : K \leq_m Y$ one can compute arbitrarily long alternating correctly labeled lists without a halting oracle. Use m simultaneously referring programs. They compute their rational queries $r(e_i)$, sort with a fixed index tie-break, and halt on alternate ranks beginning with the first. Equal adjacent queries would receive opposite labels, contradicting correctness of r. Thus all queries are distinct, with the prescribed alternating labels.

For exact Davies–Metropolis evolution and $\gamma > 0$, positive scaling keeps the energy projectors and frequency partition fixed, changes $r_\beta(\gamma\omega)$ to $r_{\gamma\beta}(\omega)$, and changes only the known final Hamiltonian unitary in addition. Hence the storage experiment for $(\gamma H, \beta)$ is equivalent to that for $(H, \gamma\beta)$ when recovery can absorb that unitary. Thermal-syndrome and thermal logical preparations correspond because their Gibbs weights agree, as do arbitrary CPTP preparations.

If a yes query x tolerates all $\gamma \in (1 - \epsilon, 1 + \epsilon)$, its nominal diagram therefore contains $((1 - \epsilon)x, (1 + \epsilon)x)$. For adjacent opposite labels $x < y$, if x stores then $y \geq (1 + \epsilon)x$. If y stores then $x \leq (1 - \epsilon)y$, giving $y/x \geq 1/(1 - \epsilon) \geq 1 + \epsilon$. Multiplication over the $m - 1$ gaps

proves the ratio bound of Theorem E. In any compact $[\beta_-, \beta_+] \subset (0, \infty)$ it permits at most

$$1 + \left\lfloor \frac{\log(\beta_+/\beta_-)}{\log(1+\epsilon)} \right\rfloor$$

ordered alternating probes.

Corollary 8.1 (Full-Pauli hot exclusion and cold probes)**.** *For a uniformly bounded finite-range qubit Hamiltonian family on a fixed-dimensional lattice, add the unit-attempt Y_i proposal to Definition 2.1. Under this exact full X, Y, Z Davies–Metropolis bath there is $\beta_h > 0$, depending only on the local bounds, such that every $0 < \beta \le \beta_h$ fails strong storage with every CPTP encoder. For a fixed strong-storage diagram and many-one halting reduction satisfying Theorem E's homogeneous rescaling promise, its computable alternating lists beginning with a yes query therefore satisfy*

$$\beta_{\max} > \beta_h (1+\epsilon)^{m-1}.$$

In particular arbitrarily long lists reach arbitrarily cold absolute temperatures. The ratio bound of Theorem E itself also holds for the X, Z bath.

Proof. We state the two additional imports precisely. Rouzé, França and Alhambra [28], Theorem II.1 of version three, give a positive volume-independent gap for their Gaussian KMS sampler using the three onsite Pauli proposals at sufficiently small inverse temperature, for fixed interaction support, incidence, norm and lattice bounds. Slezak, Scandi, França, Alhambra and Rouzé [30], Lemma 1 and Appendix A, give nondecreasing gaps as the corresponding Gaussian resolution parameter $\kappa \ge 1$ increases. Their Proposition 9 applies this monotonicity to the high-temperature $\kappa = 1$ gap $1/(2\sqrt{2}e^{1/4})$ for single-site Pauli jumps, the bound of the RFA theorem just quoted. Their frequency formula has fixed-volume limit

$$\Lambda_{\nu_1,\nu_2} \longrightarrow \delta_{\nu_1,\nu_2} e^{-(\beta\nu_1+1)^2/4}.$$

The balance drift vanishes in this limit. Fixed normalization factors are independent of volume. Finite-dimensional spectral continuity on the Gibbs-orthogonal complement of identity therefore transfers the uniform gap to the exact Gaussian-rate Davies dissipator.

For $x = \beta\omega \ge 0$,

$$e^{-(x+1)^2/4} e^x = e^{-(x-1)^2/4} \le 1;$$

for $x < 0$ the Gaussian rate is at most one. Metropolis dominates these rates pointwise. The GNS commutator-square form (3.3) compares the exact dissipative gaps. The KMS and GNS spectral gaps of a Davies dissipator agree: it preserves each modular frequency sector, and the two inner products there differ by a positive scalar. Thus the full-Pauli Metropolis gap has a positive uniform lower bound. The reference conversion in Lemma 3.1, with this gap and $-\log \sigma_{\min} = O_\beta(n)$, gives Gibbs replacement in $O_\beta(n + \log(1/\xi))$ time on arbitrary referenced inputs. This excludes every strong exponent. The alternating list begins with a yes query above β_h; apply the ratio bound already proved. □

9 What a storing endpoint can be

This section proves Theorem F. It needs no import: only that each finite cube is a computable finite optimization. We use standard notions of computable analysis [32]. A real t is *left-c.e.* if $\{q \in \mathbb{Q} : q < t\}$ is computably enumerable, and *right-c.e.* if $\{q \in \mathbb{Q} : q > t\}$ is. It is computable exactly when it is both. A subset of a rational compact interval I is *effectively*

closed if its complement in I is the union of a computably enumerable list of rational open intervals. A set is *effective* F_σ if it is the union of a uniformly effectively closed sequence.

Throughout, the interaction is computable in the sense of Section 2: the coordinate algorithm prints its Gaussian-rational Hermitian coefficient matrices exactly, and the local terms are those matrices times one common positive number c supplied by a computable name. The bath is the exact bare X, Z Davies–Metropolis bath of Definition 2.1, encoding and recovery are arbitrary CPTP maps, and the threshold is $f = 7/8$. Write S for the set of inverse temperatures $\beta > 0$ with strong storage (Definition 2.3).

9.1 A finite-cube description

For a cube size L and an integer $j \geq 1$ put $T_{L,j} = \exp(n_L^{1/j})$ and let $F_{L,j}(\beta)$ be the maximum over CPTP encoders and recoveries of the Bell overlap at time $T_{L,j}$.

Lemma 9.1 (Finite cubes are computable)**.** *On every rational compact $I \subset (0, \infty)$, each $F_{L,j}$ is computable and Lipschitz with a computable constant, uniformly in (L, j). Consequently the sets $\{\beta \in I : F_{L,j}(\beta) \geq 7/8\}$ are effectively closed, uniformly in (L, j).*

Proof. The Hamiltonian H_L is c times a Gaussian-rational Hermitian matrix H_L^0. The eigenvalues of H_L^0 are real algebraic numbers, and its spectral projectors have algebraic entries; both are computable exactly. Equality between differences of algebraic numbers is decidable. So the Bohr frequencies of H_L, which are c times those of H_L^0, fall into computable equality classes, and the jump operators q_ω of Definition 2.1 are computable and independent of β. The generator $\mathcal{L}_\beta$ depends on β only through the rates $r_\beta(\omega) = \min\{1, e^{-\beta\omega}\}$. These are computable, and on I they are Lipschitz in β with constant at most $|\omega|$. Hence $\|\mathcal{L}_\beta - \mathcal{L}_{\beta'}\|_\diamond \leq \Lambda_0 |\beta - \beta'|$ for a computable Λ_0.

Both channels are contractions, so Duhamel's formula

$$e^{T\mathcal{L}_\beta} - e^{T\mathcal{L}_{\beta'}} = \int_0^T e^{(T-s)\mathcal{L}_\beta} (\mathcal{L}_\beta - \mathcal{L}_{\beta'}) e^{s\mathcal{L}_{\beta'}}\, ds$$

gives a diamond-norm Lipschitz constant $T_{L,j}\Lambda_0$. The matrix exponential at a computable time is computable by its series with effective tails. The Bell overlap of a pair $(\mathcal{E}, \mathcal{R})$ changes by at most $\|\mathcal{E} - \mathcal{E}'\|_\diamond + \|\mathcal{R} - \mathcal{R}'\|_\diamond$ when the pair changes, and by at most the channel's diamond-norm change when β changes. The CPTP sets are compact, with Stinespring dimensions bounded by input times output dimension, and they have computable finite nets of any prescribed radius, as in Section 5. Maximizing over a net therefore approximates $F_{L,j}(\beta)$ to any precision at rational β, and $F_{L,j}$ is $T_{L,j}\Lambda_0$-Lipschitz.

For the superlevel set, enumerate rational $c \in I$ and rational $\rho, \epsilon > 0$, together with an approximation $\tilde{F}$ to $F_{L,j}(c)$ of error at most ϵ, whenever $\tilde{F} + \epsilon + T_{L,j}\Lambda_0\rho < 7/8$. Enumerate the relative open interval $I \cap (c - \rho, c + \rho)$ for each successful test. Every such interval lies in the complement, and these intervals cover every point of the complement, including any endpoint of I that it contains. □

Lemma 9.2 (Quantifier form)**.** *For every $\beta > 0$, $\beta \in S$ if and only if there are integers $j, N \geq 1$ with $F_{L,j}(\beta) \geq 7/8$ for all $L \geq N$.*

Proof. Suppose the condition holds and $L \geq N$. A maximizing pair at time $T_{L,j}$ serves every earlier time as well, because the recovery may first apply the remaining known semigroup (Definition 2.3). So $\tau(L, \beta) \geq \exp(n_L^{1/j})$, which is strong storage with $\eta = 1/j$.

Conversely, suppose $\tau(L,\beta) \geq \exp(n_L^\eta)$ for all $L \geq L_0$. Choose j with $1/j < \eta$. For $n_L \geq 2$, $\tau(L,\beta) > T_{L,j}$. By the definition of τ as a supremum, some admissible horizon exceeds $T_{L,j}$, and its encoder achieves overlap at least $7/8$ at $T_{L,j}$. Hence $F_{L,j}(\beta) \geq 7/8$. □

9.2 Proof of Theorem F

By Lemma 9.2,

$$S \cap I = \bigcup_{j,N\geq 1} K_{j,N}, \qquad K_{j,N} = \bigcap_{L\geq N} \{\beta \in I : F_{L,j}(\beta) \geq 7/8\}. \tag{9.1}$$

By Lemma 9.1, each $K_{j,N}$ is effectively closed, uniformly in (j,N): its complement is the union of all the complement enumerations for $L \geq N$. So $S \cap I$ is effective F_σ.

(i) Suppose $t \in S \cap I$ and no $\beta \in I$ with $\beta < t$ lies in S. Choose (j,N) with $t \in K = K_{j,N}$. Since $K \subset S$, K has no point of I below t, so $t = \min K$. A rational $q < \min I$ is below t. For a rational $q \in [\min I, t)$, the compact interval $[\min I, q]$ misses K, so finitely many of the enumerated complement intervals cover it. That finite cover is eventually found. No rational $q \geq t$ has such a cover, because $t \in K$. Hence $\{q : q < t\}$ is c.e. and t is left-c.e.

(ii) The symmetric argument shows that $t = \max K$, with covers of $[q, \max I]$, is right-c.e.

(iii) Let t be an isolated point of $S \cap I$. Choose a rational closed interval $J \subset I$ with $S \cap J = \{t\}$, and (j,N) with $t \in K_{j,N}$. Then $K_{j,N} \cap J = \{t\}$ is effectively closed, and t is both its minimum and its maximum. By (i) and (ii) applied inside J, t is left-c.e. and right-c.e., hence computable. □

Remark 9.3. The theorem is about each individual endpoint; the proof hard-codes a pair (j,N) and does not find it from the description. Restricting the encoders to a predetermined computable region changes nothing, and neither does the full X, Y, Z bath. The real $a + (b-a)(1-\Omega)$, with Ω Chaitin's halting probability [11], is right-c.e. and not computable, hence not left-c.e. So it can never be an attained lower storing endpoint. If instead $S \cap I = (t, \max I]$ with t in the interior of I, the storing rationals form a set that is c.e. relative to $0'$, because each $K_{j,N}$ is co-c.e.; so t is right-c.e. relative to $0'$. That is the arithmetic shadow of (9.1) used for the rational trace in Section 7. Theorem F concerns strong storage only.

10 One fixed block with a prescribed strong transition

In Theorem C temperature selects which block survives. Here the block is fixed by size alone, and temperature switches the energy barrier of that one block. Planar Ising squares sit inside the layout, and their bond pattern encodes the target temperature. Through readout cylinders they pin, or release, the checks crossed by one slab logical operator. The bath is the exact bare X, Z Davies–Metropolis bath of Definition 2.1 for the stated Hamiltonian. The memory import is the fixed-scale BDL input of Section 2.1; no mixed-scale lemma is needed. Classical spins are qubits whose terms are diagonal; an Ising bond of strength J is the term $-JZ_xZ_y$, and a field h is $-hZ_x$. Fix rationals $0 < a < b$ and a target $t \in (a,b)$ throughout.

10.1 Layered Ising squares

For a dyadic $q = m/2^k \in [0,1]$ let $D_q \subset \mathbb{Z}$ be the set of rows y for which the integer obtained by reversing the k low-order binary digits of y mod 2^k is less than m. Then D_q is 2^k-periodic with density exactly q. It does not depend on the representation of q, and $q \leq q'$ implies

$D_q \subseteq D_{q'}$: at a common level $K \geq k$, the top k digits of the reversed K-digit integer are the reversed k low digits. Appendix A checks these facts exactly.

Fix rationals $0 < J_0 < J_1$, chosen below. The *layered comparator* $P(q)$ is the Ising model on $\mathbb{Z}^2$ with every vertical bond of strength 1, horizontal bonds of strength J_1 in the rows of D_q, and J_0 in the other rows.

We use Cimasoni and Duminil-Copin [12] for positive periodic couplings on planar, locally finite, non-degenerate doubly periodic graphs, where non-degenerate means that the complementary faces are topological discs. Their Theorem 1.1 identifies the unique critical inverse temperature by equality between the high-temperature even-subgraph weight in zero mod-two homology and that in the other three classes together. Theorem 4.4(i) assumes vanishing plus magnetization throughout a neighborhood of the weight vector and gives exponential two-point decay; part (ii) gives the corresponding dual decay under positive magnetization throughout such a neighborhood. The strict phase inequalities below ensure these neighborhood hypotheses.

Lemma 10.1 (Critical layered density)**.** *Put* $\Phi_q(\beta) = 2\beta[J_0 + (J_1 - J_0)q] + \log\tanh\beta$. *The model* $P(q)$ *is subcritical at* β *if* $\Phi_q(\beta) < 0$ *and has positive spontaneous magnetization if* $\Phi_q(\beta) > 0$. *Its planar dual at* β*, with couplings* K_e^* *given by* $\tanh K_e^* = e^{-2\beta J_e}$*, is a doubly periodic layered model that is subcritical exactly when* $\Phi_q(\beta) > 0$.

Proof. Take the fundamental domain of width one and height $P = 2^k$. Each row is one vertex carrying a horizontal loop of weight $x_y = \tanh(\beta J_y)$, and consecutive rows are joined cyclically by vertical edges of weight $x_v = \tanh\beta$. In an even subgraph the vertical edges are all present or all absent, while horizontal loops are arbitrary. The class is the pair (parity of horizontal loops, presence of the vertical cycle). With $A = \prod_y (1 + x_y)$ and $B = \prod_y (1 - x_y)$, the four class sums are $(A \pm B)/2$ and $x_v^P(A \pm B)/2$. The balance of Theorem 1.1 becomes $B = x_v^P A$. Since $(1 - \tanh u)/(1 + \tanh u) = e^{-2u}$, this is $-2\beta\sum_y J_y = P\log\tanh\beta$, that is $\Phi_q(\beta) = 0$. The function Φ_q increases strictly from $-\infty$ to ∞, and couplings are monotone, so the sign of Φ_q gives the phase.

The dual of $P(q)$ has uniform horizontal couplings K_v^*, with $\tanh K_v^* = e^{-2\beta}$, and vertical couplings K_y^*, with $\tanh K_y^* = e^{-2\beta J_y}$, that depend only on the row. The same computation gives the dual balance $-2PK_v^* = \sum_y \log\tanh K_y^*$. Kramers–Wannier duality is an involution, $e^{-2K^*} = \tanh K$, so this reads $P\log\tanh\beta = -2\beta\sum_y J_y$ again. Larger dual couplings lie on the side $-2PK_v^* < \sum_y \log\tanh K_y^*$, that is, $\Phi_q(\beta) < 0$. So the dual is supercritical exactly when $P(q)$ is subcritical, and conversely. The relation $\Phi_q = 0$ is the classical periodic layered criterion of Au-Yang and McCoy [1]. Appendix A checks the class sums by enumerating even subgraphs. □

Let $g(\beta) = -\log\tanh(\beta)/(2\beta)$, which is computable, continuous and strictly decreasing on $(0, \infty)$. Choose rationals $J_0 < g(b)$ and $J_1 > g(a)$, and put $p(\beta) = (g(\beta) - J_0)/(J_1 - J_0)$. Then p is computable and strictly decreasing, $p([a, b]) \subset (0, 1)$, and $P(q)$ is subcritical at β if $q < p(\beta)$ and supercritical if $q > p(\beta)$.

A *radial square* of odd side W is the square Λ_W centered at a site 0, with vertical bonds of strength 1. Its horizontal bond from (x, y) has strength J_1 if $y \in D_{q(r)}$ and J_0 otherwise, where $r = 1 + \max(|x|, |y|)$ and $q(\cdot)$ is a computable sequence of dyadics in $[0, 1]$. A plus boundary is the field $h_x = \sum J_{xy}$ over the missing exterior bonds of each rim site; it is a sum of at most four bonds. Let $m_W(\beta)$ be the expectation of Z_0 in this square at inverse temperature β.

Lemma 10.2 (Layered squares). *1. If t is the limit of a computable sequence of rationals, there is a computable $q(\cdot)$ with $q(r) \to p(t)$. For every such $q(\cdot)$ and every $0 < \beta < t$, $m_W(\beta) \le C_\beta e^{-c_\beta W}$ for all W, with $C_\beta, c_\beta > 0$. For every $\beta > t$, $\liminf_W m_W(\beta) > 0$.*

2. If t is left-c.e., $q(\cdot)$ can be chosen computable with $q(r) \to p(t)$ and with a rational $m_* > 0$ such that $m_W(\beta) \ge m_*$ for every $\beta \ge t$ and every large W.

Proof. The schedule. From rationals $t_r \to t$, clamped to $[a, b]$, compute a dyadic $q(r)$ within 2^{-r} of $p(t_r)$; then $q(r) \to p(t)$.

Finite changes. Ferromagnetic correlations with free or plus boundary are nondecreasing in the couplings and in nonnegative fields, by Griffiths' inequalities [19]. Changing a fixed finite set E_0 of bonds changes any free two-point function, and any plus-boundary one-point function, by at most a factor κ^2 in either direction. Here κ is the product over E_0 of the larger ratio of the old and new values of $\tanh(\beta J_e)$, and the bound is uniform in the volume. To see this, expand $\langle Z_u Z_v \rangle$ in even and $\{u, v\}$-boundary subgraphs, using weights $\tanh(\beta J_e)$ and cancelling the cosh factors. Treat the plus boundary as bonds to one ghost vertex, so that the plus mean is a correlation with that vertex. Every term changes by a factor in $[\kappa^{-1}, \kappa]$.

Warm side. Let $\beta < t$. Then $p(\beta) > p(t)$; choose a dyadic $q_+ \in (p(t), \min\{1, p(\beta)\})$, or $q_+ = 1$ if $p(\beta) > 1$. Then $P(q_+)$ is subcritical at β by Lemma 10.1. Since $q(r) < q_+$ for $r \ge r_0(\beta)$, the square's couplings are at most those of $P(q_+)$ outside the bonds of radius below r_0. Raising those finitely many bonds to J_1 gives a dominating model $\hat{J}$ that differs from $P(q_+)$ on a fixed finite set. The Griffiths–Hurst–Sherman inequality [20] makes the magnetization concave in nonnegative fields. Since it vanishes at zero field,

$$m_W(\beta) \le \beta \sum_{y \in \partial\Lambda_W} h_y \langle Z_0 Z_y \rangle^{\text{free}}_{\Lambda_W, \hat{J}} \le \beta\kappa^2 \sum_{y \in \partial\Lambda_W} h_y \langle Z_0 Z_y \rangle_{P(q_+)}.$$

Free correlations increase with the volume, and the subcritical periodic model has $\langle Z_0 Z_y \rangle \le Ce^{-c|y|}$ uniformly over the finitely many period classes (Theorem 4.4 of [12]). There are $O(W)$ rim sites, each at distance at least $(W-1)/2$, which gives the bound.

Cold side. Let $\beta > t$. Choose a dyadic $q_- \in [0, p(t))$ with $q_- > p(\beta)$, or $q_- = 0$ if $p(\beta) < 0$. Then $P(q_-)$ has positive spontaneous magnetization m^* at β. Outside a finite set the square's couplings dominate $P(q_-)$; lowering those finitely many bonds gives a dominated model. Plus-boundary magnetization decreases with the volume to that of the plus state, so $m_W(\beta) \ge \kappa^{-2} m^*$ for every W large enough to contain the finite set.

Part 2. Let $\beta_j \uparrow t$ be computable increasing rationals, with $\beta_1 \ge a$, and set $q_0 = 1$. Choose a dyadic $d_j \in (p(\beta_j), \min\{1, p(\beta_j) + 2^{-j}\})$, and let $q_j = \min\{q_{j-1}, d_j\}$. Then $q_j > p(\beta_j)$ and $q_j \downarrow p(t)$, so $P(q_j)$ is supercritical at β_j, and its dual at β_j is subcritical by Lemma 10.1.

Computable dual bounds. For each j and each of the finitely many period classes u of the dual, search the boxes S of radius R about u for one with

$$\varphi(S) = \sum_{x \in S,\, y \notin S,\, xy \text{ an edge}} \tanh(K^*_{xy}) \langle Z_u Z_x \rangle^{\text{free}}_S < \tfrac{1}{2}.$$

The certified search terminates by subcritical exponential decay. Duminil-Copin and Tassion [16], Lemma 2.7, give

$$\langle Z_u Z_v \rangle^+ \le \sum_{x \in S,\, y \notin S} \tanh(K^*_{xy}) \langle Z_u Z_x \rangle^{\text{free}}_S \langle Z_y Z_v \rangle^+.$$

Their finite-volume backbone proof does not use transitivity and applies to these periodic weights. Retain one certified box per period class and let R be their maximum ℓ^∞ radius.

Reapplying the certificate for the current endpoint's class gives $2^{-\lfloor \|u-v\|_\infty/(R+1)\rfloor+1}$. Apply GHS to the interior of a box whose rim is fixed plus; the rim bonds give finite boundary fields one layer inward. Summing the free correlations there gives computable $C_j, c_j > 0$ for the wired dual one-arm bound. Full rim wiring dominates arbitrary exterior connections, and decreasing dual couplings only lowers this probability.

Radii. Choose computable increasing integers R_j, growing at least geometrically, with $C_j e^{-c_j r} \le e^{-2L^{1/4}}$ whenever $L \ge R_j/20$, $r \ge L^{1/3}/100$ and L exceeds an absolute constant. Put $q(r) = q_j$ for $r \in [R_j, R_{j+1})$, and $q(r) = q_1$ for $r < R_2$. The masks are nested. So inside radius $R < R_{j+1}$ the couplings dominate $P(q_j)$, and since $\beta_j \le t$, the dual at any $\beta \ge t$ is dominated there by the dual of $P(q_j)$ at β_j. Part 1 applies because $q(r) \to p(t)$.

Magnetization at t. Work in an infinite-volume free random-cluster measure at $\beta \ge t$, a limit of finite-volume ones; it is the Edwards–Sokal representation [17]. For $D = 2^s$, consider two events: an open circuit in the annulus between radii D and $2D$, and an open crossing of a corridor of width comparable to D from radius D to radius $4D$. Failure of either forces a closed dual crossing of spatial diameter at least cD, contained with its local arm boxes inside radius $8D$, by planar duality [6]. That path contains a dual site with a dual arm of radius $D^{1/3}$ in its own box. By the radii condition and a union bound over $O(D^2)$ sites, failure at scale D has probability at most $CD^2 e^{-cD^{1/4}}$. Choose a computable D_0 beyond R_2 for which the sum over $D \ge D_0$ is less than $1/2$. On the success of all scales the circuits form an infinite cluster.

Also require every bond in the ball of radius $3D_0$ to be open. Each has conditional probability at least $\tanh(\beta J_{\min}) \ge \tanh(aJ_{\min})$, with $J_{\min} = \min\{J_0, 1\}$. All these events are increasing, so the FKG inequality gives

$$\mathbb{P}(0 \leftrightarrow \infty) \ge \tfrac{1}{2} \tanh(aJ_{\min})^{N_0},$$

where N_0 is the number of forced bonds. Let ϕ^{f}_β be the infinite-volume free random-cluster measure just used. Wire the exterior neighbors of Λ_W together, retaining the bonds from Λ_W to those neighbors at their original couplings. This gives exactly the fixed-exterior plus boundary of Section 10.1. Boundary-condition monotonicity and the Edwards–Sokal coupling give

$$m_W(\beta) \ge \phi^{\mathrm{f}}_\beta(0 \leftrightarrow \Lambda_W^{\mathrm{c}}) \ge \phi^{\mathrm{f}}_\beta(0 \leftrightarrow \infty) \ge \tfrac{1}{2} \tanh(aJ_{\min})^{N_0}.$$

The bound holds for every $\beta \ge t$. Choose a positive rational m_* strictly below its last expression. □

10.2 The readout cylinder

Fix integers $K, J_r \ge 1$ and an odd W, and put $M = W^2$. The *readout cylinder* is the square lattice $\mathbb{Z}_{2W} \times \{0, \dots, (M+3)W\}$ with free ends and Ising bonds of strength J_r. At height 0 there are at most W distinct *ports*, each with a field of absolute value at most K. Above height $2W$ the cylinder is cut into consecutive bands of height W. In the middle of the ith band, $1 \le i \le M$, a *source* site b_i carries a field $h \ge 0$. There are no other fields.

Lemma 10.3 (Readout cylinder)**.** *Fix $a > 0$ and $h_0 > 0$ and an integer $K \ge 1$. There is an integer J_r such that the following holds for every $\beta \ge a$. If the source field satisfies $\beta h \ge d > 0$ for all large W, then every port p has*

$$\Pr(Z_p = -1) \le \epsilon_\beta + o_{\beta,d}(1) \qquad (W \to \infty), \qquad \epsilon_\beta \le e^{-2\beta K - h_0}/4.$$

The $o(1)$ term is uniform over the port fields and over p.

Proof. A plus circuit without the port fields. First drop the port fields. In the Edwards–Sokal representation [17], add a ghost vertex with plus spin joined to each source. Given all other edges, an edge of coupling u is open with probability at least $\tanh u$. So sequential exposure couples the lattice edges above independent bond percolation of density $\tanh(\beta J_r)$, and the ghost edges above independent ones of density $\tanh d$. Put $v = 1 - \tanh(\beta J_r) \le 2e^{-2\beta J_r}$ and assume $3v < 1/2$.

Consider the band between heights W and $2W$ and the M source bands. Require that each band contains an open circuit winding around the cylinder, and that each union of two consecutive bands contains an open path from its bottom to its top. Then all these circuits lie in one cluster. Failure of a winding circuit or a crossing produces a closed dual path of at least $W - 2$ edges. Counting a starting edge and at most three choices per step, the failure probability is at most $C(M+1)W^2(3v)^{W/2}$.

A source b_i fails to reach both rims of its band within the band only if a closed dual circuit or rim-to-rim arc separates it. For a separator of length $\ell < W/4$ there are at most $C(\ell+1)^2 3^\ell$ choices; longer separators cost $CW^2(3v)^{W/4-2}$. So a source is bad with probability at most $1/8$ once J_r is large in units of $1/a$. Distinct bands use disjoint edges, so more than half the sources are bad with probability at most $2^M 8^{-M/2}$. If at least half the sources reach both rims, at least $M/2$ independent ghost edges have their lattice end in the crossing cluster, and all are closed with probability at most $(1 - \tanh d)^{M/2}$.

If one is open, the circuit in the band $[W, 2W]$ has plus spin. Restoring the port fields changes the probability of any event by at most $e^{2\beta KW}$. So the probability that no plus winding circuit exists in that band is at most

$$e^{2\beta KW}\left[C(M+1)W^2(3v)^{W/2} + 2^{-M/2} + (1-\tanh d)^{M/2}\right],$$

which tends to zero, uniformly in the port fields, once $(2\beta J_r - \log 6)/2 > 2\beta K + 1$; here $M = W^2$ matters.

The local minus cost. Now work with the port fields. On the event that $Z_p = -1$ and a plus winding circuit exists in the band $[W, 2W]$, draw the interfaces through disagreeing bonds, rounded at faces by a fixed rule. Some interface separates p from the circuit: either an arc with both ends on the bottom rim, or a winding loop below the circuit. Let A be the region it cuts off on the side of p. It lies below height $2W$, so it contains no source. If the interface crosses ℓ bonds, A meets at most 3ℓ bottom sites, because each bottom column of A leaves A through a distinct crossed bond. Flipping all spins of A is injective for a given interface and side. It gains $2\beta J_r \ell$ in bond energy and loses at most $6\beta K\ell$ in port fields. So that interface has probability at most $e^{-(2\beta J_r - 6\beta K)\ell}$.

There are at most $C(\ell+1)^2 3^\ell$ such interfaces: a short arc starts within distance $\ell + 2$ of p, and a long one starts anywhere in the $O(W^2) \le C(\ell+1)^2$ positions of the band. Hence

$$\Pr(Z_p = -1) \le o(1) + C\sum_{\ell \ge 1} (\ell+1)^2 \left[3e^{-2\beta J_r + 6\beta K}\right]^\ell.$$

For $\beta \ge a$ one integer J_r makes the sum at most $e^{-2\beta K - h_0}/4$ and also meets the conditions above. □

10.3 A barrier comparison after summing classical registers

Lemma 10.4 (Barrier comparison)**.** *Take n_d data qubits with bounded-support commuting Pauli checks spanning a stabilizer group of rank $n_d - 1$, and a Hermitian Pauli U that commutes*

with every check and represents a nontrivial logical. Write U as a product of l single-qubit X or Z factors on distinct sites. Add n_f classical qubits and put $n_0 = n_d + n_f$. Let $H = E(s,z)$ be a sum of uniformly bounded terms, each a diagonal function of a bounded number of check signs s_c and classical Z's, with bounded incidence. Suppose that, given s, the Gibbs law of z factors over disjoint groups of at most b spins.

Write $w(s) = \sum_z e^{-\beta E(s,z)}$ on formal sign assignments, and for a set V of factors of U let s^V be s with the signs of the checks that anticommute with $\prod_{v\in V} u_v$ reversed. Assume that for every prefix V in some order of the factors and every s,

$$e^{-B} \le w(s^V)/w(s) \le e^{B}.$$

Then, under the exact bare X, Z Davies–Metropolis bath at β, the channel at every time

$$t \ge \mathrm{poly}_\beta(n_0)\, e^{B + C_\beta b}\,[n_0 + \log(1/\xi) + 1]$$

is within trace norm ξ of an entanglement-breaking channel on every input with reference.

Proof. The Dirichlet form. By (3.3), with b_1 bounding the number of Bohr pieces of a proposal and $r_* > 0$ the smallest rate,

$$\mathcal{E}(F) := -\Re\langle F, \mathcal{L}_\beta^* F\rangle_\sigma \ge \frac{r_*}{2b_1} \sum_q \|[q,F]\|_\sigma^2\,, \tag{10.1}$$

where q runs over all X_i, Z_i. The Hamiltonian part is skew.

Sectors. The $n_d - 1$ independent checks, U, and the n_f classical Z's are n_0 independent commuting Paulis. Each commutes with H and maps every proposal to plus or minus itself, so its conjugation commutes with the dynamics and preserves σ. The joint character sectors are therefore invariant and orthogonal. The trivial sector consists of the operators diagonal in the joint eigenbasis, a rank-one basis; projecting onto it is the complete pinching $\mathcal{D}$, which commutes with the channel.

Suppose some local check or classical Z, call it g, has character -1. Then $\|[g,F]\|_\sigma^2 = 4\,\|F\|_\sigma^2$. Writing g as a product of $O(1)$ proposals and telescoping, conjugation through each prefix changes σ by a bounded factor, because a bounded number of bounded terms change. So $\|F\|_\sigma^2 \le C \sum_q \|[q,F]\|_\sigma^2$.

The remaining sector. It is positive on every check and classical Z and negative on U. Use the data eigenbasis $|x\rangle$, $x = (s,j)$, with j the sign of U. An operator F in this sector maps $|x,z\rangle$ to $f(x,z)|x^*,z\rangle$, where x^* flips j, and $\|F\|_\sigma^2 = \mathbb{E}_\pi |f|^2$ for $\pi(x,z) \propto e^{-\beta E(s,z)}$. Let π_x be the conditional law of z, $g(x) = \mathbb{E}_{\pi_x} f(x,\cdot)$, G the operator with entries g, $\bar\pi(x) \propto w(s)$, and $V_f = \mathbb{E}_{\bar\pi} \mathrm{Var}_{\pi_x} f$. Then $\|F\|_\sigma^2 = \|G\|_{\bar\pi}^2 + V_f$.

Each group of at most b classical spins has a ratio of largest to smallest conditional weight at most $e^{C_\beta b}$. Comparing with uniform independent flips, whose Walsh characters have gap 2, and tensorizing over groups gives

$$V_f \le e^{C_\beta b} \sum_{\text{classical } i} \|[X_i,F]\|_\sigma^2\,. \tag{10.2}$$

For a data proposal q, $q|x\rangle$ is a phase times $|qx\rangle$, and the ratios π_x/π_{qx} and $\bar\pi(x)/\bar\pi(qx)$ lie in $[C'^{-1}, C']$, because only a bounded number of bounded terms change. Write $g(x) - \chi g(qx)$, with χ the phase, as the π_x-mean of $f(x,\cdot) - \chi f(qx,\cdot)$ plus χ times $\mathbb{E}_{\pi_x} f(qx,\cdot) - \mathbb{E}_{\pi_{qx}} f(qx,\cdot)$. Jensen's inequality bounds the first part by the commutator of F. The second is at most $C' \mathrm{Var}_{\pi_{qx}} f(qx,\cdot)$ in square. Summing gives

$$\|[q,G]\|_{\bar\pi}^2 \le 2\,\|[q,F]\|_\sigma^2 + C'' V_f. \tag{10.3}$$

Write $U = q_l \cdots q_1$ in the given order. Telescope $[U, G]$ along this path, conjugate the diagonal weight $\bar{\pi}$ through each prefix, and apply Cauchy–Schwarz. Then

$$4\,\|G\|_{\bar{\pi}}^2 = \|[U, G]\|_{\bar{\pi}}^2 \le l\, e^B \sum_q \|[q, G]\|_{\bar{\pi}}^2 . \tag{10.4}$$

Equations (10.1)–(10.4) give $\|F\|_\sigma^2 \le \mathrm{poly}_\beta(n_0) e^{B + C_\beta b} \mathcal{E}(F)$ on every nontrivial sector. So the generator has gap $\lambda \ge [\mathrm{poly}_\beta(n_0) e^{B + C_\beta b}]^{-1}$ off the trivial sector.

Operational conclusion. Write $N_t = \mathcal{D} N_t + (\mathrm{id} - \mathcal{D}) N_t$. The first term is entanglement-breaking, since $\mathcal{D}$ measures a complete rank-one basis. For a test A with $\|A\| \le 1$, $\|(\mathrm{id} - \mathcal{D}) N_t^* A\|_\sigma \le e^{-\lambda t}$. The conversion of Lemma 3.1, with $\sigma_{\min} \ge 2^{-n_0} e^{-C_\beta n_0}$, bounds the referenced error by $2^{n_0} \sigma_{\min}^{-1/2} e^{-\lambda t}$. □

10.4 Thermal witnesses for coupled syndrome laws

Lemma 10.5 (Coupled witness bound)**.** *Fix the BDL seed and scale ℓ_0, let $h_0 = b_{\mathrm{suff}}(\ell_0)$ from Lemma 6.1 with $\ell_f = \ell_0$, and let $2k \ge 4$. Let μ be a law on the realizable syndromes of C_{2k}, in both sectors, of the form $\mu(\sigma) \propto w(\sigma)$, where $w > 0$ is defined on all formal sign assignments. Suppose that changing any one check from satisfied to violated multiplies w by at most e^{-h_0}, uniformly in all other coordinates. Then $\mu(B_k) \le 4\, 2^{-2^{k-1}}$.*

Proof. Let τ be a witness, realizable in its sector by BDL Lemma 5.8, and let $\sigma \supseteq \tau$ be realizable. Remove the checks of τ one at a time; this gives $w(\sigma) \le e^{-h_0|\tau|} w(\sigma + \tau)$. The map $\sigma \mapsto \sigma + \tau$ is injective into realizable syndromes, so $\mu(\tau \subseteq \sigma) \le e^{-h_0|\tau|}$. The counting part of the proof of Lemma 6.1, at $\ell_f = \ell_0$, uses the law only through this inequality. It gives $2\, 2^{-2^{k-1}}$ per sector. □

10.5 The construction

Constants. Let h_0 be as in Lemma 10.5. Choose an integer K with $2aK > h_0 + \log 4$. Choose J_r by Lemma 10.3 for a, h_0, K, and J_0, J_1 as in Section 10.1. Every coupling is one of finitely many fixed rationals.

The block and its slab logical. Take the fixed-scale level-$2k$ BDL block with its computable bounded-density embedding (Section 2.1). Its qubit count and width are at most $A_1 B_1^k$. Cut space into slabs normal to z whose constant thickness exceeds the check range, and let $\mathcal{S}_A$ be every other slab and $\mathcal{S}_B$ the rest.

For a CSS code with one logical qubit, either some X-type logical is supported on $\mathcal{S}_A$ or some Z-type logical is supported on $\mathcal{S}_B$. Indeed, a Z-type logical can be moved off $\mathcal{S}_A$ by multiplying with Z-checks exactly when it commutes with every X-type logical supported on $\mathcal{S}_A$. So the Z-type logicals supported on $\mathcal{S}_B$ form the annihilator of the X-type logicals supported on $\mathcal{S}_A$, and the two dimensions add to one. This is the cleaning argument of Bravyi and Terhal [9]; Appendix A checks the count exactly on random small codes.

No check meets two slabs of the same color. So such a logical splits into a product of single-slab operators of the same type, each commuting with every check, and one of them is a nontrivial logical. Compute one by finite linear algebra with a fixed tie rule. Call it $U = \prod_v u_v$: a product of only X's or only Z's, in one slab.

Decoration. For a check S_c, let D_c be the set of sites where u_v anticommutes with the factor of S_c. It has even size, since U commutes with S_c. If $D_c = \emptyset$ the check gets the

term $K(I - S_c)$. Otherwise choose an anchor site in D_c and give the check the contact term $K(I - S_c Z_p)$ with a port p of a readout cylinder.

Classify the decorated checks by the y coordinate of the anchor and a finite *type*: the offset pattern of D_c, the anchor's depth in the slab, and a multiplicity index. Then each (x, y, type) cell holds at most one anchor. Along each (y, type) line, cut the x positions into intervals of $W = 2k^2 + 1$ consecutive positions. Each interval that holds an anchor gets one readout cylinder.

The cylinder is folded into a double row of W cells along x in its own slot, and extends in z over height $(M + 3)W$. Its ports sit on the forward row at the anchors' x positions, at the height of the slab. Each source b_i is joined by an Ising bond of strength 1 to the center of a private radial square of side W (Lemma 10.2). The square carries the plus boundary fields and occupies an extra slot over its band. More explicitly, for an interval beginning at x_0, map circular indices $j = 0, \ldots, 2W - 1$ to $x_0 + j$ for $j < W$ and to $x_0 + 2W - 1 - j$ for $j \geq W$, using two distinct slots. Consecutive circular sites are then adjacent cells or distinct slots of one cell. Choose every source on the forward row at the middle x position of its interval. Put its square in the x, z plane, centered at the source, in a third slot and within that source's band. Since W is odd, the square occupies exactly the interval's W columns. Source bands and neighboring intervals have disjoint interiors; a fixed endpoint convention makes their site sets disjoint. The number of types is fixed, so these three slots per type, together with the data slots, give a uniform occupancy bound. Contact range is at most the slab thickness plus the check range and a fixed rounding allowance. Each cylinder and its squares use $O(W^4)$ spins and height $O(W^3)$.

Different cylinders and squares use disjoint spins. Contact terms have bounded range, and only a bounded number of slots per cell is used. All terms commute, and the only operators on the data are the original checks times classical signs. The decorated block has active size and width at most $A_2 B_2^k$, with computable A_2, B_2.

Shells. Choose a rational $\alpha > 0$ with $B_2^{\alpha s} \leq 2^{s/3}$ for all sufficiently large s, and a computable s_0 beyond which the decorated blocks, with their fixed margins, fit in the indicated shells. In each x-shell $[2^s, 2^{s+1})$ with $s \geq s_0$ install one decorated block of depth $2\lfloor \alpha s \rfloor$, with $k = \lfloor \alpha s \rfloor$ and $W = 2k^2 + 1$. Its other coordinates lie below 2^s. Every unused slot, including those in the earlier shells, is a free qubit. For $2^s \leq L < 2^{s+1}$ with $s \geq s_0 + 1$, let B_L be the block in shell $s - 1$, which is complete in Λ_L; choose B_L arbitrarily for the remaining finitely many sizes. A coordinate algorithm computes every local term by a finite computation, using the schedule $q(\cdot)$ of Lemma 10.2.

10.6 Proof of Theorem G

The single-bond field. Sum out a square at a fixed sign of its source. Expanding $e^{\beta Z_c Z_b} = \cosh \beta \, (1 + Z_c Z_b \tanh \beta)$ shows that the square acts on Z_b as a field whose exponent is exactly $\operatorname{atanh}(\tanh(\beta) m_W(\beta))$, where m_W is the square's own plus-boundary center mean.

Cold side. Let $\beta > t$, or $\beta \geq t$ in the left-c.e. case. By Lemma 10.2, $m_W(\beta) \geq m' > 0$ for large W, so each source has exponent at least $d = \operatorname{atanh}(\tanh(a) m') > 0$. Fix check signs s and let $w(s)$ be the product of the bare factors and of the cylinder partition functions with port fields $K s_c$. Making a satisfied decorated check violated multiplies w by the cylinder expectation of $e^{-2\beta K Z_p}$. By Lemma 10.3 this is at most

$$e^{-2\beta K} + e^{2\beta K}(\epsilon_\beta + o(1)) \leq e^{-2aK} + e^{-h_0}/4 + o(1) \leq e^{-h_0}$$

for large k, uniformly in the other coordinates. A bare check gives the factor $e^{-2\beta K} \le e^{-h_0}$. Lemma 10.5 then bounds the unstable probability of the joint Gibbs law by $4\,2^{-2^{k-1}}$.

The thermal logical encoding of B_L at β prepares the joint Gibbs law of the syndrome, cylinder and square spins of that block in its decoder frame, and Gibbs spectators. Lemma 2.4, with these auxiliary registers, gives Bell overlap at least $1 - 8A_1B_1^kT2^{-2^{k-1}}$ at every time T. This tends to one at $T = e^{c2^k}$ for $0 < c < (\log 2)/2$. With $k \ge \alpha(s-1) - 1$ and $n = q_0L^3 \le q_0 2^{3(s+1)}$, every $0 < \eta < \alpha/3$ works at all large sizes. The block location depends only on L.

Warm side. Let $\beta < t$ and $L = 2^{s+1}$; every installed block is complete. By Lemma 10.2, $m_W(\beta) \le C_\beta e^{-c_\beta W}$, so each source exponent is at most that.

Order the factors of U by y, then x, then the remaining coordinates. After a prefix V, a check is reversed exactly when $|D_c \cap V|$ is odd. Only anchors within the check range of the current y can be reversed. For such a y and a type, anchors far to the left of the current x share one parity, fixed by the type and the completed rows, and anchors far to the right share another. Hence, apart from $O(1)$ cylinders near the current site, each cylinder has all of its contacts reversed or none, and bare checks are never reversed.

Reversing all ports of a cylinder and flipping all its spins changes its partition function by a factor in $[e^{-2M\bar{h}}, e^{2M\bar{h}}]$, through the sources only, where $\bar{h} \le C_\beta e^{-c_\beta W}$ is the source exponent. A partial cylinder costs at most $e^{2\beta KW}$. With at most $A_2B_2^k$ cylinders,

$$|\log w(s^V)/w(s)| \le C\beta KW + 2A_2B_2^kW^2C_\beta e^{-c_\beta W} = O_\beta(1+W),$$

since $W = 2k^2 + 1$. Given the data signs, the classical spins split into groups of $b = O(W^4)$: one cylinder with its squares. Lemma 10.4 therefore makes each large block, with error $\xi = n^{-2}$, close to entanglement-breaking by time $\exp(O_\beta(k^8))\,\mathrm{poly}(n) = \exp(O_\beta((\log n)^8))$.

The finitely many small early blocks obey Lemma 3.1, and free qubits depolarize. The channel is the tensor product of its component channels (Section 2). Replacing components one at a time adds the errors, so the whole channel is within $o(1)$ of an entanglement-breaking channel on every input with reference. That stays true at all later times, because an entanglement-breaking channel followed by any channel is entanglement-breaking. Every Bell overlap is then at most $1/2 + o(1)$, which defeats $7/8$ and every $\eta > 0$ along these sizes.

Endpoint. For left-c.e. t, part 2 of Lemma 10.2 supplies the cold estimate at every $\beta \ge t$, and part 1 the warm one. Chaitin's Ω is left-c.e. [11], and so is $a + (b-a)\Omega$. □

10.7 Endpoint classes and the common-encoder bound

Corollary 10.6 (Closed and robust endpoints)**.** *1. Let I be a rational compact interval and t an interior point. For any interaction covered by Theorem F, if $S \cap I = [t, \max I]$ then t is left-c.e. Conversely, for every left-c.e. $t \in (a,b)$ the layout of Theorem G has $S \cap [a,b] = [t,b]$.*

2. Suppose a computable interaction as in Theorem F has a computable increasing sequence of sizes L_s, rationals $1/2 \le f_{\mathrm{lo}} < f_{\mathrm{hi}} \le 1$ and a rational $\eta > 0$, such that at every rational $\beta > t$ the optimal Bell overlap at time $\exp(n_{L_s}^\eta)$ is eventually at least f_{hi}, and at every rational $\beta < t$ eventually at most f_{lo}. Then t is limit-computable. Conversely, every limit-computable $t \in (a,b)$ arises in this way from Theorem G, with $L_s = 2^{s+1}$, $f_{\mathrm{hi}} = 3/4$ and $f_{\mathrm{lo}} = 5/8$.

Proof. Part 1 is Theorem F(i) together with Theorem G. For the first half of part 2, fix a rational $\theta \in (f_{\mathrm{lo}}, f_{\mathrm{hi}})$. For rational β let the bit $e_s(\beta)$ record whether an approximation, with error $1/s$, of the finite optimum at size L_s exceeds θ; it is computable by Lemma 9.1. With the halting oracle one can search for a stage after which $e_s(u)$ no longer changes. Dovetail this search over rational pairs $u < v$, and accept a pair when the bits at u and v stabilize to 0 and 1. An accepted pair brackets t, and arbitrarily narrow pairs are accepted. So t is computable from the halting set, which is to say limit-computable.

For the converse, Theorem G gives overlap tending to one at every $\beta > t$ at all large sizes, with one η. It gives overlap at most $1/2 + o(1)$ at every $\beta < t$ on the sizes 2^{s+1}. □

Remark 10.7 (Why Theorem D does not apply). Theorem D concerns m temperatures at each of which a different register A_i survives while the rest of the layout is replaced. These registers are disjoint. One encoder serving all m temperatures must therefore spread one qubit over m disjoint registers, and monogamy limits the worst overlap to $(m+1)/(2m)$. In Theorem G the register that survives at every storing temperature is the same block B_L. In that argument $m = 1$, and the bound reads 1. Temperature does not choose a register; it switches the energy barrier of one register. Theorem D remains true for the layout of Theorem C.

What still depends on temperature is the preparation of classical labels: the encoder samples the syndrome, cylinder and square spins of the fixed block from their Gibbs law at the actual β. These labels carry no logical information. A preparation independent of temperature is not established here.

Remark 10.8 (Further endpoint classes). Remark 9.3 shows that an open storing ray $(t, \max I]$ has an endpoint that is right-c.e. relative to $0'$. A longer construction, not reproduced here, realizes every such endpoint with a fixed-block finite-palette layout. It also realizes every effective F_σ subset of a rational compact interval as an exact storing set. We state neither as a theorem. In those constructions, failure at temperatures outside the storing set is shown only along sparse sequences of complete sizes. At some of those temperatures the same layout stores along infinitely many other sizes. Their extra strength over Corollary 10.6 comes from the quantifier "at every sufficiently large size" in Definition 2.3, as in Theorem C.

Remark 10.9 (Scope). Theorem G is a statement about a layout: the fabrication rule computes the bond masks of the squares from the program for t, and it computes the decoration. The radii R_j of Lemma 10.2 are computable but astronomically large. The theorem says nothing at $\beta = t$ unless t is left-c.e. It uses exact frequency resolution, as every result here does.

11 Bounded-width commuting-Pauli chains

Rapid thermalization of translation-invariant one-dimensional commuting Hamiltonians is established by Bardet, Capel, Gao, Lucia, Pérez-García and Rouzé [3]. The quantum-information bound below uses local erasure and dephasing instead of full equilibrium mixing.

Theorem I (One-dimensional quantum loss)**.** *Consider any family of one-dimensional commuting-Pauli Hamiltonians with uniformly bounded interaction range, qubits per slice and single-proposal energy change J, under the exact bare X, Z Davies–Metropolis bath. For every finite β and integer $k \geq 2$, every syndrome-diagonal encoder, with arbitrary CPTP recovery and reference, has Bell overlap at most $1/2 + 1/(2k)$ after a size-independent time. This includes thermal-syndrome and pure joint-syndrome ground-space encodings, open chains, cycles and fixed-width tubes. If $H = \alpha \sum_c (1 + 2^{-j(c)-2})(I - S_c)/2$ for the commuting Pauli*

checks S_c, a gain $\alpha > 0$ and a proper coloring j of the check intersection graph, as in (5.1)*, the same conclusion holds for every CPTP encoder. Consequently weak and strong storage fail in the respective preparation classes.*

Proof. **A bounded-range diagonalizing Clifford.** Group consecutive slices into cells at least as long as the check diameter and at most a fixed multiple of it. Enlarge the last cell when needed. Every check is onsite or meets one neighboring pair. A cycle with fewer than four cells, and any residual short chain, has bounded total size and can be one cell.

At cell C, let L_C, R_C be the binary Pauli subspaces spanned by the restrictions of left-edge and right-edge checks. Their cross symplectic pairings vanish, since such checks meet only in C and commute globally. Onsite checks commute with both and with each other. In L_C split off successive symplectic planes: choose an anticommuting pair, subtract its pairings from the remaining vectors, and continue. The remaining radical is orthogonal to L_C. The whole of R_C and all onsite checks are orthogonal to the removed planes. Repeat for R_C. The two radicals and the onsite subspace in the remaining symplectic complement are isotropic. Extend their basis to a symplectic basis and make those vectors single-qubit Z coordinates. A finite symplectic basis transformation is implemented by a Clifford, using Hadamard, phase and CNOT elimination on one Pauli pair at a time. Each cell has bounded size.

A check on edge $C, C+1$ now acts on dedicated right factors of C, dedicated left factors of $C+1$, and the remaining classical Z coordinates. Its dedicated restrictions commute with all other dedicated restrictions on that edge: the classical factors commute and the original checks commute. A Clifford on those two dedicated factors simultaneously diagonalizes their commuting list. Different edges have disjoint dedicated factors, so these changes form a bounded-depth circuit. All checks are now products of virtual single-qubit Z operators, with signs retained. The resulting Clifford U and its inverse carry a single-site Pauli to a uniformly bounded-radius Pauli string. This is the Pauli form of the interaction-algebra decomposition of Bravyi and Vyalyi [10]; the symplectic argument above supplies it directly.

The exact prepared algebra. An input diagonal in the joint physical check labels stays in that algebra under every Bohr proposal. The Hamiltonian drift vanishes there. Let $a = e^{-\beta J} > 0$. On this invariant algebra the exact proposal generator has the CPTP extension consisting of unconditional Pauli conjugations at rate a and residual attempts at rate $1 - a$, with Kraus operators

$$q\sqrt{(r_q - a)/(1-a)}, \qquad \sqrt{(1-r_q)/(1-a)}.$$

For $a = 1$ omit the residual attempts. On each source sector these maps have precisely the acceptance rate r_q. Conjugate by U. The base maps are complete dephasings by bounded Pauli strings at rate $2a$. Residual Kraus operators are Pauli strings times diagonal local matrices or diagonal local matrices. Every conditional update therefore commutes as a channel with each virtual Z_i-dephasing. All base string dephasings commute as channels, even for anticommuting Pauli strings.

For virtual qubit i, express $U^\dagger X_i U$ and $U^\dagger Z_i U$ as products of physical single-site X, Z generators, using both for each Y factor. Their lists have uniformly bounded size, radius and incidences. The corresponding virtual strings generate X_i, Z_i. Let A_i contain their union of virtual supports. In an initial unit time bin, require all listed base dephasings and silence of every residual whose support meets A_i. Then these dephasings move to the bin's beginning. Their product is the Pauli twirl over their generated subgroup, which absorbs both X_i- and

Z_i-dephasing and hence full replacement at i. This is an operator identity with arbitrary reference. It permits simultaneous insertion of replacements at all certified targets.

Let D bound base and residual source/target incidences after enlarging lists if necessary. Split base clocks into D independent subclocks, assigning one to each requesting target. Represent residual silence by the OR construction of Lemma 3.2. With private padding this gives iid sufficient erasure marks of probability

$$p_E = (1 - e^{-2a/D})^D [1 - (1 - e^{-(1-a)})^{1/D}]^D > 0.$$

Use identical independent constructions in m later bins. Every later mark permits virtual Z_i-dephasing at that bin's input. Covariance of every conditional update moves it to the original input. Initial replacements commute with these dephasings. Thus over $m + 1$ bins the extended channel is postprocessing of an iid flagged channel with erasure probability $p = p_E$, intact probability $u = (1 - p_E)^{m+1}$ and classical-measurement probability $1 - p - u$. Conditional clock sampling constructs this postprocessing explicitly and is independent of the input. Choose m so $p \geq (k-1)u$, and apply Lemma 3.3. Conjugating the input and output by the fixed Clifford changes no encoder or reference bound. Equality with the physical dynamics on syndrome-diagonal inputs proves the first assertion.

All coherent inputs in the resolved subclass. For the colored coefficient list, equation (5.2) holds on every input at fixed β, with $\zeta > 0$ depending only on its bounded local energy change. Remove the known unitary at readout. Divide time into two-unit bins. In each bin's second half, assigned subclocks certify complete pinches of all required individual check labels, exactly as in Lemma 5.1; their probabilities are positive for every fixed finite β. These pinches commute to the first half's input. Every conditional map also commutes with each virtual Z_i-dephasing, since the frequency masks are diagonal and the proposals are Pauli strings in virtual coordinates.

For target i, require all labels needed by every early center whose support meets A_i. At those centers use the exact base/residual realization on the pinched algebra; every unprepared center misses A_i. Require silence of the overlapping residuals and occurrences of the assigned base lists. Splitting clocks and applying Lemma 3.2 to both the pinch and silence source families gives iid sufficient marks of one fixed positive probability across targets and bins. No lower bound of $1/2$ is needed. The early string twirl absorbs full replacement, which commutes with the inserted stabilizer pinches because they are Pauli twirls. Later marks certify virtual Z_i-pinches movable to time zero by the same covariance. The preceding erasure/measurement/intact comparison and broadcast bound therefore hold on arbitrary inputs. This proves the resolved-subclass assertion. □

12 Discussion

The address construction makes weak passive storage undecidable with a fixed finite alphabet. Its equal-spectrum version isolates the role of the physical bath operators: a global Clifford preserves energy levels while changing which local errors connect their eigenspaces. The fixed-material gain construction then turns the inverse temperature into the only query, with a relativized degree statement that retains the halting oracle in the join.

Strong storage requires more than arbitrarily large protected blocks. Their size and spacing must be compared with the volume of every specimen. In Theorem C thermometer arms and programmed omissions realize this comparison, and the encoder chooses a completed block;

for that layout the common-preparation bound measures the cost of the choice. Theorem G removes the choice. One block, fixed by size, has a barrier switched by temperature, and its transition can be any limit-computable temperature. Theorem F marks the limit of this freedom: for every interaction it covers, a lowest storing inverse temperature in an interval, if storage holds there, is left-c.e.

The constructions are existence results with astronomically large constants. The weak gain material has infinitely many computable coupling values. Both strong layouts have finite microscopic palettes but computable, nonperiodic placement rules: Theorem C through infinitely many species and an omission schedule, Theorem G through computed bond masks in its Ising squares. The Σ_2^0 complexity of Theorem C reflects the every-large-size quantifier through its omission schedule. Storage is a property of the Hamiltonian, the bath's local coupling operators and the preparation class together; by Theorem A, equal spectra do not determine it.

A strong finite-palette layout with one noncomputable transition now exists. Whether one material, periodic or generated by a fixed finite substitution, can have a noncomputable strong transition remains open. For the present construction, the missing step is a fixed rule for the joint registration of the code and the devices that set its barrier. One possible encoding mechanism is ruled out: on a nonempty effectively closed shift, a local frequency common to every invariant measure is computable, by a compactness search over the invariant measures. This observation does not exclude other geometric or thermodynamic mechanisms for a noncomputable transition. Thermal computation is one possible route; no necessity statement is established here. Two further questions are open. One is a preparation independent of temperature, in particular pure ground-state startup for the imported blocks under the bare bath. The other is temperature hardness for a finite-resolution bath with general static perturbations.

How this work was done

This work used a combination of AI models, including Claude Opus 5.5, GPT-6.1 Sol and GPT-6 Astra, directed by the authors. Every theorem has a complete written proof here. Independent refutation-first examinations checked the research proofs behind each result against the imported sources. Exact integer and rational computations verify the finite identities that the proofs use (Appendix A). The positive three-dimensional memory is imported from version 1 of BDL, Sections 3–5.

A Finite computations and reproducibility

The proofs above are analytic. The only finite code property used in amplification is the commutation, independence and distinct single-error syndromes of the four displayed five-qubit generators. The program `verify_finite.py` checks it by exact binary symplectic arithmetic. It also checks the following finite identities independently of the asymptotic proofs.

Check	Method and scope
Five-qubit code	Binary ranks, pairwise symplectic products, all fifteen single-Pauli error syndromes.
Amplification	Exact rational value of $f(p)/p$ at $p = 63/500$ and a rational positive lower bound on its derivative.
Thermometer exponents	Rational identities for all endpoint exponents; an analytic sufficient middle-interval starting value using $e^{-x} \le 2/x^2$.
Isospectral blocks	Exact signed-Pauli projector traces versus direct classical parity enumeration for all eight cuts of a three-data-site CSS list, with a spare species.
Colored signs	Exact dyadic signed-sum distinctness for palettes of sizes one through twelve. The general proof is in Section 5.
Weighted support	Exhaustive binary subspaces up to four coordinates with a specified finite rational weight grid and all coordinate subsets. The general Fourier proof is Lemma 3.4.
Cloning	Exact rational sums for the occupation-state marginal formula for $1 \le m \le 32$.
Bit-reversal masks	Exact density of $D_{m/2^k}$ over one period, independence of the representation, and nesting $D_q \subseteq D_{q'}$ for all dyadics with $k \le 6$ (Section 10.1).
CSS cleaning count	For 60 seeded random CSS codes on at most seven qubits with $k \ge 1$ and every region A, the X-type logical dimension on A plus the Z-type logical dimension on its complement equals k (Section 10.5).
Layered torus classes	Enumeration of all even subgraphs of the width-one torus with $P \le 5$ rows and random rational weights; the class balance equals $\prod(1 - x_y) - \prod x_v \prod(1 + x_y)$ (Lemma 10.1).

The program uses only the Python 3.10 or later standard library, with exact integers and fractions, and runs in a few seconds. From the directory that contains it, run

```
python verify_finite.py
```

It prints `PASS` followed by the name of each of the ten groups in the table, then `ALL FINITE CHECKS PASSED`; any failed identity stops it with an assertion error. The paper itself builds from `main.tex` with `pdflatex` and `bibtex`.

These computations verify finite algebraic facts and sufficient inequalities, all in exact arithmetic. They do not simulate the large physical systems, certify an asymptotic lifetime by sampling, decide a noncomputable set, or replace the imported BDL statements.

The finite five-qubit certificate

The four syndrome bits use the generator order displayed in Section 5. Every nonzero four-bit word occurs once in the table. Any two generators have exactly two anticommuting overlaps, so they commute. Their X supports are the edges of the path 3–1–4–2–5. A nonempty subset of these edges has a leaf, whose X bit cannot cancel; hence the generators are independent. The common plus eigenspace therefore has dimension two. The sixteen listed syndrome sectors, including the zero syndrome for identity, are disjoint. Measurement of the syndrome followed by the inverse listed single-error Pauli recovers the original logical state, including its reference.

Site	X	Y	Z
1	0001	1011	1010
2	1000	1101	0101
3	1100	1110	0010
4	0110	1111	1001
5	0011	0111	0100

B Imported results and their scope

All positive three-dimensional results use BDL [2], version one. Section 2.1 states the fixed-scale memory, embedding-existence and decoder-frame imports: Theorem 3.16, Lemmas 3.3 and 3.15, Definition 5.1, Lemma 5.2, Corollary 5.3, Theorem 5.13 and Corollary 5.14. Lemma 6.1 identifies the additional single-step cohomology, coarse-graining and witness imports by number and proves their mixed-scale composition. Its use of Lemma 3.14, Claim 5.5, Lemmas 5.6 and 5.8–5.12 and Eqs. (54), (218)–(235) is explicit there.

The full-Pauli high-temperature calibration corollary additionally imports the two Gaussian-sampler results stated in its proof, from [28, 30]. Theorem G uses only the fixed-scale BDL statements, with the counting of Lemma 6.1 at a single scale. It imports Theorems 1.1 and 4.4 of Cimasoni and Duminil-Copin [12], stated in Section 10.1. It also uses standard tools stated where they are applied: Griffiths' and the Griffiths–Hurst–Sherman inequalities [19, 20], the Edwards–Sokal coupling [17], planar duality [6], and the correlation inequality of Duminil-Copin and Tassion [16]. The other cited works provide context and attribution. The channel convergence, finite separability, sufficient-mark coupling, monogamy inequality, automata reduction, finite-resource threshold, thermometer estimates, small-field comparison, common-encoder construction, bounded-query diagonalization and bounded-width dynamical argument are written out above. The five-qubit code and symmetric cloner are specified by their matrices or generators, so their needed properties can be verified directly.

References

[1] Helen Au-Yang and Barry M. McCoy. Theory of layered Ising models: Thermodynamics. *Physical Review B*, 10:886–891, 1974. doi:10.1103/PhysRevB.10.886.

[2] Shankar Balasubramanian, Margarita Davydova, and Ting-Chun Lin. A passive self-correcting quantum memory in three dimensions. arXiv:2605.10943v1, 2026. Version 1; Sections 3–5.

[3] Ivan Bardet, Ángela Capel, Li Gao, Angelo Lucia, David Pérez-García, and Cambyse Rouzé. Rapid thermalization of spin chain commuting hamiltonians. *Physical Review Letters*, 130:060401, 2023. doi:10.1103/PhysRevLett.130.060401; arXiv:2112.00593.

[4] Johannes Bausch, Toby S. Cubitt, Angelo Lucia, and David Pérez-García. Undecidability of the spectral gap in one dimension. *Physical Review X*, 10:031038, 2020. doi:10.1103/PhysRevX.10.031038; arXiv:1810.01858.

[5] Johannes Bausch, Toby S. Cubitt, and James D. Watson. Uncomputability of phase diagrams. *Nature Communications*, 12:452, 2021. doi:10.1038/s41467-020-20504-6; arXiv:1910.01631.

[6] Vincent Beffara and Hugo Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. arXiv:1006.5073, 2012.

[7] Richard Beigel. *Query-limited reducibilities*. PhD thesis, Stanford University, 1987. Report STAN-CS-88-1221; electronic thesis archive.

[8] Richard Beigel, William I. Gasarch, John Gill, and James C. Owings. Terse, superterse, and verbose sets. *Information and Computation*, 103:68–85, 1993. doi:10.1006/inco.1993.1014.

[9] Sergey Bravyi and Barbara Terhal. A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes. *New Journal of Physics*, 11: 043029, 2009. doi:10.1088/1367-2630/11/4/043029; arXiv:0810.1983.

[10] Sergey Bravyi and Michael Vyalyi. Commutative version of the k-local hamiltonian problem and common eigenspace problem. *Quantum Information and Computation*, 5: 187–215, 2005. arXiv:quant-ph/0308021.

[11] Gregory J. Chaitin. A theory of program size formally identical to information theory. *Journal of the ACM*, 22:329–340, 1975. doi:10.1145/321892.321894.

[12] David Cimasoni and Hugo Duminil-Copin. The critical temperature for the Ising model on planar doubly periodic graphs. *Electronic Journal of Probability*, 18:no. 44, 1–18, 2013. Theorems 1.1 and 4.4; doi:10.1214/EJP.v18-2352; arXiv:1209.0951.

[13] Toby S. Cubitt, David Pérez-García, and Michael M. Wolf. Undecidability of the spectral gap. *Nature*, 528:207–211, 2015. doi:10.1038/nature16059; arXiv:1502.04135.

[14] E. B. Davies. Markovian master equations. *Communications in Mathematical Physics*, 39:91–110, 1974. doi:10.1007/BF01608389.

[15] Seth Douglas and Nidhal Mghirbi. Undecidable quantum storage: no algorithm can tell from the noise rate whether a fixed circuit remembers. arXiv:2611.29134v1, 2026. Version 3.1.0; archive concept record doi:10.5281/zenodo.24254845.

[16] Hugo Duminil-Copin and Vincent Tassion. A new proof of the sharpness of the phase transition for Bernoulli percolation and the Ising model. *Communications in Mathematical Physics*, 343:725–745, 2016. doi:10.1007/s00220-015-2480-z; arXiv:1502.03050.

[17] Robert G. Edwards and Alan D. Sokal. Generalization of the Fortuin–Kasteleyn–Swendsen–Wang representation and Monte Carlo algorithm. *Physical Review D*, 38:2009–2012, 1988. doi:10.1103/PhysRevD.38.2009.

[18] Nicolas Gisin and Serge Massar. Optimal quantum cloning machines. *Physical Review Letters*, 79:2153–2156, 1997. doi:10.1103/PhysRevLett.79.2153; arXiv:quant-ph/9705046.

[19] Robert B. Griffiths. Correlations in Ising ferromagnets. II. external magnetic fields. *Journal of Mathematical Physics*, 8:484–489, 1967. doi:10.1063/1.1705220.

[20] Robert B. Griffiths, C. A. Hurst, and S. Sherman. Concavity of magnetization of an Ising ferromagnet in a positive external field. *Journal of Mathematical Physics*, 11: 790–795, 1970. doi:10.1063/1.1665211.

[21] Patrick M. Hayden. Robust passive quantum memory in three dimensions, part i: Thermal and zero-temperature perturbative stability. arXiv:2611.21749v1, 2026. Version 1.

[22] Markus Holzer, Martin Kutrib, and Andreas Malcher. Multi-head finite automata: Characterizations, concepts and open problems. arXiv:0906.3051, 2009.

[23] Carl G. Jockusch, Jr. Semirecursive sets and positive reducibility. *Transactions of the American Mathematical Society*, 131:420–436, 1968. doi:10.1090/S0002-9947-1968-0220595-7.

[24] Alastair Kay, Dagomir Kaszlikowski, and Ravishankar Ramanathan. Optimal cloning and singlet monogamy. *Physical Review Letters*, 103:050501, 2009. doi:10.1103/PhysRevLett.103.050501; arXiv:0901.3626.

[25] Martin Kummer. A proof of beigel's cardinality conjecture. *Journal of Symbolic Logic*, 57:677–681, 1992. doi:10.2307/2275299.

[26] Raymond Laflamme, Cesar Miquel, Juan Pablo Paz, and Wojciech Hubert Zurek. Perfect quantum error correction code. *Physical Review Letters*, 77:198–201, 1996. doi:10.1103/PhysRevLett.77.198; arXiv:quant-ph/9602019.

[27] James Purcell, Zhi Li, and Toby Cubitt. Chaitin phase transition. arXiv:2410.02600, 2024.

[28] Cambyse Rouzé, Daniel Stilck França, and Álvaro M. Alhambra. Efficient thermalization and universal quantum computing with quantum gibbs samplers. arXiv:2403.12691v3, 2026. Version 3, Theorem II.1.

[29] Naoto Shiraishi and Keiji Matsumoto. Undecidability in quantum thermalization. *Nature Communications*, 12:5084, 2021. doi:10.1038/s41467-021-25053-0; arXiv:2012.13889.

[30] Samuel Slezak, Matteo Scandi, Daniel Stilck França, Álvaro M. Alhambra, and Cambyse Rouzé. Polynomial-time thermalization and gibbs sampling from system-bath couplings. arXiv:2601.16154v1, 2026. Version 1, Lemma 1 and Appendix A.

[31] Kristan Temme. Thermalization time bounds for Pauli stabilizer Hamiltonians. arXiv:1412.2858, 2014.

[32] Klaus Weihrauch. *Computable Analysis: An Introduction*. Springer, Berlin, 2000. doi:10.1007/978-3-642-56999-9.

[33] R. F. Werner. Optimal cloning of pure states. *Physical Review A*, 58:1827–1832, 1998. doi:10.1103/PhysRevA.58.1827; arXiv:quant-ph/9804001.

[34] Andrew C. Yuan. One-shot any code. arXiv:2610.02137v1, 2026. Version 1.