

Channel-resolved memory time of Haah's cubic code under reversible rate filtering and temporally correlated baths

Hikaru Wakaura

hikaruwakaura@gmail.com

QuantScape Inc

Research Article

Keywords: Self-correcting quantum memories , Haah's cubic code.

Posted Date: October 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11241681/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Channel-resolved memory time of Haah’s cubic code under reversible rate filtering and temporally correlated baths

Hikaru Wakaura^{1,*}

¹*QIRI (Quantum Integrated Research Institute Inc.), Tokyo 107-0061, Japan*

(Dated: October 2, 2026)

Self-correcting quantum memories would store logical information without a real-time decoder, and Haah’s cubic code is the best studied three-dimensional candidate, but how its memory time depends on the environment beyond a flat Markovian bath has not been mapped. Here we simulate the thermal syndrome dynamics of the cubic code at $\beta = 4.5$ –5 for up to 13718 qubits under detailed-balance rate filters on each of its three energy channels, under an Ornstein–Uhlenbeck modulation that gives the bath a correlation time, and at rescaled code energies. The zero-energy diffusion channel controls the memory time: suppressing it by 10^3 lengthens the memory time 80- to 120-fold for $L = 7$ –13, the gain enters as a prefactor rather than a steeper size law, and it outlives the flat-bath optimum, where the memory time peaks between $L = 13$ and 17 and falls by 39% at $L = 19$ while the filtered bath still grows to $L = 17$ and gains a factor 45 at $L = 19$. The pair-creation channel is invisible over three decades; a bath correlation time adds a factor of two to eleven, consistent with annealed and quenched hazard averages; and a uniform 12% energy rescaling raises the size exponent from 2.7 to 4.3 along the Bravyi–Haah line, which no gapped pseudomode we solve provides. These results form a channel-resolved finite-size benchmark for bath engineering of a passive three-dimensional memory; they establish neither asymptotic scaling nor an advantage over active decoding.

I. INTRODUCTION

Storing a quantum state for longer than the coherence time of its physical carriers is the first task any quantum technology must solve, and it is solved today by active error correction: stabilizers are measured, a classical decoder infers the most likely error, and a correction is applied before the next round. Surface-code memories on superconducting processors now operate below threshold, with the logical error rate suppressed by $\Lambda = 2.14 \pm 0.02$ each time the code distance grows by two and a real-time decoder with $63 \mu\text{s}$ average latency (distance five) at a $1.1 \mu\text{s}$ cycle [1]; low-density parity-check codes promise the same performance with roughly ten times fewer qubits [2], and the resource cost of the decoder and of the universal gate set on top of it is a central concern [3]. In every such scheme the decoder is part of the physics: its latency, its classical throughput, and rare correlated error events, which limit repetition-code performance in the same device [1], set the overhead of the memory.

A self-correcting memory removes the decoder from the storage phase altogether [4–7]. If the Hamiltonian that defines the code confines thermally created defects, the environment itself performs the correction and syndrome measurement is needed only once, at readout. For two-dimensional stabilizer codes with local generators the energy barrier between logical states is bounded by a constant independent of the size [8], so they cannot be self-correcting by this criterion [7, 9], and the relaxation time of the two-dimensional toric code is bounded independently of its size [10]; in three dimensions the cubic

code of Haah [11], which has no string-like logical operators, is the best studied candidate: under a Markovian Davies bath its memory time grows as a power of the linear size L whose exponent is proportional to the inverse temperature β , $T_{\text{mem}} \simeq L^{2.93\beta-10.52}$, up to an optimal size $L^* \simeq e^{0.78\beta-0.87}$ beyond which entropy wins [12, 13], so that at any fixed nonzero temperature the memory time is bounded in L [14]. Very recently a three-dimensional stabilizer Hamiltonian with exponential memory time was constructed [15], so passive three-dimensional memories exist at least in principle. The question of which three-dimensional codes are at least partially self-correcting has been reopened: layer codes have a linear energy barrier, the optimal scaling in three dimensions, and a memory time proven to grow exponentially in L up to a size that itself grows with β [16, 17]; new constructions from product codes have been proposed with the explicit aim of suppressing the entropic factors that cap the cubic code at its optimal size [18, 19]; and a sufficient condition for self-correction has been formulated at the level of finite-temperature phases, under which a Gibbs state connected by a locally reversible channel circuit to a zero-temperature code state is a memory with stretched-exponential lifetime under any quasi-local Lindbladian that fixes it [20]. The cubic code is not known to satisfy that condition, and in kinetic Monte Carlo of random layer codes with a cluster decoder the memory time is likewise found to follow a power law up to a temperature-dependent peak [17], so the finite-size regime below the optimal size is, for now, the regime in which the memory time of a three-dimensional code is actually measured, and it is the regime studied here.

What the environment can do for such a memory has been asked mostly for the two-dimensional toric code. Chesi, Röthlisberger and Loss showed that a super-

* h.wakaura@qiri.co.jp

Ohmic bath, whose thermal rate vanishes at zero frequency, removes direct anyon hopping and, in a toric code supplemented with long-range repulsion, improves the size scaling of the lifetime [21]; the lattice-geometry dependence in an Ohmic bath was studied by kinetic Monte Carlo [22], and the finite-size crossover of the toric-code relaxation time has been characterized [23]. Coupling the code to cavity modes or to a lattice of massless bosons induces effective interactions and an anyon creation penalty that grows with L [24, 25] (within a perturbative range of sizes in the cavity case [24]), but this divergence requires exactly massless bosons and is destroyed by generic local perturbations that gap them [14]; a bosonic bath also generates spatially and temporally correlated errors that can break active correction [26]. A separate, non-equilibrium line uses *designed* dissipation or measurement-free control to protect a code space [27], to trap [28] or translate [29] defects, to run a cellular-automaton decoder continuously [30], or, as a dissipative quantum cellular automaton with a proven threshold in two dimensions, to store and compute autonomously [31]; in the classification of passive correction as a dissipative phase, detailed balance is argued not to be the essential ingredient [32], and dynamical decoupling remains the reference active technique against non-Markovian noise at the single-qubit level [33, 34]. On the open-systems side, non-Markovian dynamics is rigorously classified and measurable [35–38], and a structured bath with matching two-time correlations can be represented by damped pseudomodes under a Lindblad equation [39], such Markovian embeddings being non-unique, since several auxiliary-mode schemes arise as different unravelings of the same bath self-energy [40]. However, for a three-dimensional code no study has separated, on the same footing, the effect of the spectral weight at each allowed energy change, of the bath correlation time, and of a renormalization of the code energies, on the memory time and on its size dependence, and without that separation one cannot say which bath properties a passive memory should be engineered for.

In this work we provide that separation for the cubic code as a finite-size numerical benchmark. We treat the Davies rates of the three allowed energy changes as independent reversible filters, add an externally modulated zero-energy rate with a tunable correlation time, rescale the code energies uniformly, and, in a one-defect pseudomode model, ask what a locally damped structured mode does to the defect population. We find that on the plateau of pair-creation rates the zero-energy channel has by far the largest effect on the lifetime, with the barrier channel weakly relevant and the creation channel taking over only below a creation-limited crossover, that a bath correlation time between the cluster lifetime and the memory time adds a factor of about two described by an annealed hazard average, and that a uniform energy rescaling moves the size exponent along the Bravyi–Haah line. The result shows which levers a bath engineer has at fixed temperature and code, and what remains to be

computed before any of them can be promoted to an asymptotic or operational statement.

II. MODEL AND METHODS

A. Code and thermal dynamics

We use Haah’s cubic code (Code 1 of Ref. [11]) on an $L \times L \times L$ periodic lattice with two qubits per site, $N = 2L^3$, and stabilizers G_c^X, G_c^Z on every elementary cube c , built from the corner table of Ref. [12]. The Hamiltonian is $H = -J \sum_c (G_c^X + G_c^Z)$ with $J = 1/2$. Energy units used throughout: one violated stabilizer (a defect) costs 1; a single-qubit error flips four stabilizers and costs $\Delta = 4$; the barrier channel changes the energy by ± 2 . We verified the construction by checking that all stabilizers commute and that the number of logical qubits follows the exact formula $k(L) = 2[1 - 2q_2 + 2^{r+1}(q_2 + 12q_{15} + 60q_{63})]$ of Ref. [11] ($q_n = 1$ if $n \mid L$, $2^r \parallel L$) for $L = 2-7$ and 15. All simulations use odd L not divisible by 15, for which $k = 2$ and the logical operators are known.

Thermal noise is modeled by a Davies generator [41] with single-qubit X jump operators; by the X – Z duality of the code the Z sector behaves identically, and simultaneous (Y -type) errors are not modeled. Because the Hamiltonian is a sum of commuting stabilizers, the dynamics in this sector reduces exactly to a classical continuous-time Markov chain on error configurations [12, 13]; we track the error frame e itself, not only its syndrome, so that logical errors are decided by the anticommutation of e plus correction with a logical operator. Flipping qubit q toggles the four Z stabilizers containing it, changing the energy by $\Delta E = 4 - 2n_{\text{viol}}(q)$, where $n_{\text{viol}}(q) \in \{0, \dots, 4\}$ is the number of already violated stabilizers among them. We use Metropolis rates $\gamma(\Delta E) = \gamma_0 \min(1, e^{-\beta \Delta E})$ [42, 43] and simulate with the rejection-free Bortz–Kalos–Lebowitz algorithm [44], the protocol of Ref. [12]. Time is measured in units of $1/\gamma_0$.

B. Three phenomenological environmental knobs

Reversible rate filters. A Davies generator sees a bath only through the rates at the energy changes the code allows, $|\Delta E| \in \{4, 2, 0\}$. We multiply the rates of each channel by a factor s_4, s_2, s_0 (pair creation and annihilation, barrier crossing, and zero-energy hopping), scaling forward and backward processes together so that detailed balance and the Gibbs distribution over error configurations are preserved exactly. The flat Metropolis bath is $(1, 1, 1)$. We use these filters as a phenomenological model of a reversible bath with a structured response; we do not construct the microscopic bath that realizes a given triple, and we note that a thermal bosonic rate

near zero frequency behaves as $J(\omega)/(\beta\omega)$, so a vanishing $J(0)$ (Ohmic) does not by itself give $s_0 = 0$; only a super-Ohmic rate limit does [21].

Temporal correlations. To give the bath a correlation time without changing its mean rates or the Gibbs distribution, we let the diffusion rate fluctuate, $s_0(t) = \exp[\xi(t) - \sigma^2/2]$, with ξ an Ornstein–Uhlenbeck process [45] of correlation time τ_c and stationary variance σ^2 , so that $\langle s_0 \rangle = 1$; $\xi(0)$ is drawn from the stationary distribution. Rates are held piecewise constant over $\tau_c/20$ while zero-energy moves exist and ξ is advanced with its exact transition kernel otherwise. This is an external modulation shared by all sites, not a derived thermal reservoir.

One-defect pseudomode model. A qubit with gap $\Delta = 4$ (the energy of one error) is coupled to one damped mode, in the spirit of pseudomode embeddings [39]: $H = \frac{\Delta}{2}\sigma^z + \omega_0 a^\dagger a + g(\sigma^+ a + \sigma^- a^\dagger)$, with the mode damped at rate κ by the local thermal Lindblad pair $\sqrt{\kappa(n+1)}a$, $\sqrt{\kappa n}a^\dagger$, $n = (e^{\beta\omega_0} - 1)^{-1}$. The qubit then sees a Lorentzian spectral density centered at ω_0 with width κ and memory time $1/\kappa$; we use the model as a minimal structured bath and do not claim that it is the exact embedding, in the sense of the equivalence theorem of Ref. [39] (Gaussian bosonic environment, matching two-time correlation functions), of a specific continuous thermal bath, and such embeddings are in any case not unique [40]. We solve the Lindblad equation exactly (dimension 2×13 ; the archived cutoff sweep from 8 to 28 shows no change at the 10^{-6} level; the two-defect model uses 11 oscillator levels and is converged at the 10^{-12} level) and compute separately the exact reduced Gibbs state of qubit plus mode. Its bare-defect population is the trace observable $p_{\text{def}} = \text{Tr}[\Pi_e e^{-\beta H}/Z]$ with $\Pi_e = |e\rangle\langle e| \otimes I$, which is basis-independent and continuous through resonance.

C. Readout, memory time, and fit protocol

At readout we apply the renormalization-group decoder of Ref. [12]: defects are grouped into 2^p -connected clusters for $p = 0, 1, \dots$; each cluster is tested for neutrality by solving, over $\mathbb{F}_2$, for an X operator supported on all lattice sites of the cubes in the cluster’s minimal enclosing box, which is the 1-neighborhood of the box in the sense of Definition 1 of Ref. [12], where a neutral cluster is shown to be annihilated by an operator with that support; neutral clusters are erased and the procedure is repeated up to the code’s topological-order scale $L_{\text{tqo}} = L - 2$. Our neutrality test is the generic linear-algebra version, not the broom algorithm of Ref. [12]. Following that reference, a trial decoding of a copy of the syndrome is performed every $T_{\text{ec}} = e^{4\beta}/100$, and the memory time of a run is the first trial at which the decoder either fails to clear all defects or returns a logical error. The observable is therefore the first-failure time of this specific decoder; a failed neutrality test does not by itself prove that no recovery exists. For the fixed-

rate series the lifetimes are consistent with an exponential distribution (coefficient of variation 0.97 ± 0.09 ; the modulated-rate point at $\tau_c = 10^{10}$ has 1.72), and we report the sample mean T_{mem} with its standard error over $n = 100$ runs ($n = 60$ at $\beta = 5$). All filtered series use the same random-number seeds as the flat bath, so differences between series are paired; because per-run lifetimes are not archived the covariance between series is unknown, and cross-series differences are quoted with quadrature-combined errors whose direction of bias is not established. The Ornstein–Uhlenbeck series use a different seed sequence and are not paired with the flat bath.

Size exponents are obtained by weighted least squares of $\ln T_{\text{mem}}$ on $\ln L$ with weights $(T_{\text{mem}}/\text{SEM})^2$ on the common window $L = 5\text{--}11$, which excludes the point at which the flat bath saturates; we report χ^2/dof , the error scaled by $\sqrt{\max(1, \chi^2/\text{dof})}$, and the window systematic $\Delta p_{\text{win}} = |p(5\text{--}11) - p(7\text{--}13)|/2$. A series is assigned an exponent only if $\chi^2/\text{dof} < 10$ on its window (a declared, conventional cutoff; the strongest diffusion filter, with $\chi^2/\text{dof} = 38$, is not). Two exponents are called consistent if they differ by less than a declared tolerance of 0.5, which corresponds to a 4% change of β in the Bravyi–Haah form. As an internal size diagnostic we also quote

$$\Lambda_{\text{passive}}(L) \equiv T_{\text{mem}}(L+2)/T_{\text{mem}}(L), \quad (1)$$

with errors from both numerator and denominator. The Bortz–Kalos–Lebowitz kernel is compiled with numba and runs are distributed over 16 cores; the kernel was validated against an exact uniformization solution of the master equation on all 2^{16} error configurations at $L = 2$ (Appendix A) and against the pure-Python implementation of the same decoder, with which it shares the code construction and neutrality test and from which only the kinetic inner loop differs ($1.24(8) \times 10^8$ with $n = 200$ against $1.20(17) \times 10^8$ with $n = 40$ at $L = 9$, $\beta = 4.5$).

III. RESULTS

A. Markovian baseline

Figure 1 shows the memory time in a flat bath at $\beta = 4.5$ and 5.0. On the common window the exponents are 2.73 ± 0.17 ($\chi^2/\text{dof} = 1.2$, window systematic 0.37) at $\beta = 4.5$ and 4.10 ± 0.22 ($\chi^2/\text{dof} = 0.7$) at $\beta = 5.0$, against $2.93\beta - 10.52 = 2.67$ and 4.13 from the fit of Ref. [12] over $L \leq 33$; at $\beta = 4.5$ the growth saturates at $L = 13$, consistent with $L^* = 14.0$. Two sizes beyond the optimum confirm the turnover directly: $T_{\text{mem}} = 2.21(20) \times 10^8$ at $L = 17$, equal to the $L = 13$ value $2.26(22) \times 10^8$ (ratio 0.98 ± 0.13), and $1.35(15) \times 10^8$ at $L = 19$ (ratio 0.61 ± 0.09 to $L = 17$), so the memory time has a maximum between $L = 13$ and 17 and decreases beyond it. A peak of the same kind, at a size that grows with β , is seen in kinetic Monte Carlo of random layer codes with a cluster decoder [17]. These two sizes are not used




FIG. 1. **The Markovian memory time of the cubic code follows the Bravyi–Haah power law up to an optimal size set by entropy.** Memory time versus linear size L in a flat Metropolis bath at $\beta = 4.5$ (blue, $n = 100$ runs per point) and $\beta = 5.0$ (red, $n = 60$). The $\beta = 4.5$ points at $L = 17$ and 19 lie beyond the optimal size and are not used in fits. Dashed lines: $L^{2.93\beta-10.52}$ normalized at $L = 5$ and drawn up to $L = 14$; dotted vertical lines: $L^* = e^{0.78\beta-0.87}$. Error bars are standard errors of the mean.

in any fit; all 200 of their runs end in a decoder failure, with 34 and 50 defects on average at failure against 7.5 at $L = 9$. In 499 of 500 runs the failure is the decoder’s inability to annihilate a residual cluster and in one run a logical error, as in Ref. [12]; the former is a decoder-failure count, not a proof that the residual clusters are intrinsically uncorrectable. The flat-bath size diagnostic is $\Lambda_{\text{passive}} = 1.69 \pm 0.21$ for $L = 9 \rightarrow 11$ and 1.12 ± 0.14 for $L = 11 \rightarrow 13$ at $\beta = 4.5$, decreasing toward unity as L^* is approached, as any finite polynomial growth must.

B. Which rate channel controls the memory time

Figure 2 and Table I decompose the bath into its three channels at $\beta = 4.5$. Over the tested range $s_4 \in [0.01, 10]$ the pair-creation channel has no measurable effect: at $L = 9$ the memory time is $1.14, 0.96, 1.20, 1.16 \times 10^8$ for $s_4 = 0.01, 0.1, 1, 10$, and the same holds at every $L \geq 7$ (the $L = 5$ point at $s_4 = 0.01$ is twice the flat value, a small-lattice effect). This cannot hold for arbitrarily small s_4 : no failure can precede the first error, so $T_{\text{mem}} \geq e^{4\beta}/(2L^3\gamma_0 s_4)$, and the memory becomes creation-limited below $s_4^* = 1/(Ne^{-4\beta}T_{\text{mem}}^{\text{flat}}) = 3.8 \times 10^{-4}$ at $L = 9$. Two points below the plateau confirm this: at $s_4 = 3 \times 10^{-4}$ and 10^{-4} the memory time rises to $2.89(26) \times 10^8$ and $5.56(57) \times 10^8$, in agreement with the additive form $T_{\text{mem}}^{\text{flat}} + T_{\text{first}}(s_4) = 2.7 \times 10^8$ and 5.7×10^8 with $T_{\text{first}} = e^{4\beta}/(2L^3\gamma_0 s_4)$ (Fig. 2b). Above s_4^* the plateau has a simple reading: scaling creation and annihilation together leaves the equilibrium abundance of

four-defect clusters at its Boltzmann value, and failure is caused by clusters that separate and become charged, which proceeds through the other two channels. The barrier channel is weakly relevant, with logarithmic sensitivity $d \ln T_{\text{mem}}/d \ln s_2 = -0.29$ at $L = 9$ (-0.26 at $L = 13$; three-point fits over $s \in \{0.1, 1, 10\}$). The zero-energy diffusion channel dominates, with sensitivity -0.60 (-0.66 at $L = 13$), and it has not saturated: $s_0 = 0.1, 0.01$, and 0.001 give 4.5, 30, and 103 times the flat-bath memory time at $L = 9$ (the last value from a single size).

Within the declared tolerance, the size exponent of every series that a power law describes differs from the flat bath by less than 0.5 on matched windows: $s_2 = 0.1$ and $s_0 = 0.1$ differ from the flat bath by 0.21 and 0.02, and on the 7–13 window the two s_4 series differ from it by 0.01 and 0.26. This is a statement about the fitted finite-size exponent on a stated window, with a window systematic (0.37 for the flat bath, twice its statistical error; 0.36 for $s_0 = 0.1$, comparable to its statistical error of 0.38); it is not a determination of an asymptotic exponent, which would require data on both sides of L^* for each filter. For $s_0 = 0.01$ and $s_0 = 0.001$ the growth from $L = 7$ to 13 is weaker than in the flat bath (on the 7–13 window, where a power law is acceptable, the exponents are 0.92 ± 0.23 and 1.19 ± 0.19 against 1.98 ± 0.28), and the gain over the flat bath at $s_0 = 0.001$ is $40 \pm 5, 121 \pm 17, 103 \pm 13, 80 \pm 11, 78 \pm 10$ for $L = 5, 7, 9, 11, 13$, so the gain from freezing diffusion appears, on the present data, as a large prefactor rather than as steeper size dependence. The prefactor outlives the flat-bath optimum: at $s_0 = 0.01$ the memory time is $6.49(54) \times 10^9$ at $L = 17$ and $6.05(65) \times 10^9$ at $L = 19$ (ratios 1.27 ± 0.16 for $13 \rightarrow 17$ and 0.93 ± 0.13 for $17 \rightarrow 19$; $n = 100$, flat-bath seeds; all runs end in decoder failure with 11 and 14 defects on average), where the flat bath has already fallen by 39%, so the gain grows from 22.5 ± 3.0 at $L = 13$ to 29 ± 4 at $L = 17$ and 45 ± 7 at $L = 19$ (Fig. 2a). Suppressing diffusion therefore moves the size at which the memory time peaks outward, from between $L = 13$ and 17 to $L \geq 17$; the two sizes are too few to fit the shifted optimum. This differs from the interacting toric code of Ref. [21], where a super-Ohmic bath improves the scaling because the anyon gap itself grows with L ; in the cubic code the gap and the logarithmic energy barrier are fixed by the code.

C. A bath correlation time adds a hazard average with annealed and quenched limits

Figure 3 shows the effect of the correlation time. With $\sigma = 1$ the memory time changes by at most 20% for τ_c between 10^{-2} and 10^6 hop times ($0.97 \pm 0.13, 1.13 \pm 0.14, 1.20 \pm 0.16, 0.92 \pm 0.12, 1.17 \pm 0.16$ relative to the flat bath, errors including the flat-bath denominator). With $\sigma = 2$ it grows to 1.60 ± 0.21 at $\tau_c = 1$ and 2.12 ± 0.30 and 1.97 ± 0.26 at $\tau_c = 10^2$ and 10^6 ; the difference between $\tau_c = 1$ and 10^2 is 2.0σ (numerator errors only, since both

TABLE I. **Among the plateau scans the zero-energy channel has the largest sensitivity, and no series that a power law describes has a size exponent outside the declared tolerance on its stated window.** Memory time at $\beta = 4.5$ ($n = 100$; the $s_0 = 0.01$ series extends to $L = 17$ and 19 in Fig. 2a), the exponent on $L = 5$ – 11 with scaled error and χ^2/dof , the window systematic, and the three-point sensitivity at $L = 9$. For s_4 the fit excludes the anomalous $L = 5$ point (window 7–13). The $s_0 = 0.01$ series is not described by a single power law ($\chi^2/\text{dof} = 38$ on 5–11: a $9.5\times$ jump from $L = 5$ to 7 followed by growth of only $1.8\times$ from $L = 7$ to 13 , against $3.2\times$ for the flat bath) and is given no exponent; the $s_0 = 0.001$ series behaves in the same way ($\chi^2/\text{dof} = 41$; $9.4\times$ from $L = 5$ to 7 , then $2.1\times$ from $L = 7$ to 13).

Filter	$T_{\text{mem}}(9)/10^8$	$T_{\text{mem}}(13)/10^8$	exponent	$\chi^2/\text{dof}; \Delta p_{\text{win}}$	$\frac{d \ln T_{\text{mem}}}{d \ln s}$
flat	1.20(11)	2.26(22)	2.73(17)	1.2; 0.37	–
$s_4 = 0.01$	1.14(12)	2.66(26)	1.97(22) ^a	1.2; –	+0.04
$s_4 = 10$	1.16(10)	2.77(27)	2.24(23) ^a	1.3; –	
$s_2 = 0.1$	2.62(26)	5.49(53)	2.51(18)	1.0; 0.07	–0.29
$s_2 = 10$	0.69(7)	1.63(19)	–		
$s_0 = 0.1$	5.37(50)	12.6(12)	2.70(38)	5.4; 0.36	–0.60
$s_0 = 0.01$	35.5(36)	51.0(47)	none	38; 1.20	
$s_0 = 0.001$	124(11)	176(15)	none	41; 1.32	
$s_0 = 10$	0.35(3)	0.61(6)	–		

^awindow $L = 7$ – 13 ; the flat bath on the same window gives 1.98(28).




FIG. 2. **Reversible rate filtering acts through the zero-energy and barrier channels on the plateau $s_4 \in [0.01, 10]$; below the creation-limited crossover the pair-creation channel takes over.** (a) Memory time versus L at $\beta = 4.5$ for each channel filter (s_4, s_2, s_0). (b) Memory time at $L = 9$ versus the filter factor of a single channel, with the other two at unity; the dotted line is the naive kinetic scaling $\propto 1/s$, the dashed line the flat-bath value, and the thin solid line the additive creation-limited form $T_{\text{mem}}^{\text{flat}} + T_{\text{first}}(s_4)$, with the crossover $s_4^* = 3.8 \times 10^{-4}$ marked; the two leftmost s_4 points were added to test it. Error bars are standard errors over $n = 100$ runs.

share the flat-bath denominator), so the data locate the onset of the effect at or below one hop time and only marginally resolve a crossover.

The correlation times $\tau_c = 1$ – 10^6 satisfy $\tau_{\text{cluster}} \lesssim \tau_c \ll T_{\text{mem}}$: a defect cluster lives for of order one to a hundred hop times, while $T_{\text{mem}} \sim 10^8$, so the bath is renewed many times before a failure (the point $\tau_c = 10^{-2}$ lies below the cluster lifetime and is included only to bound the fast limit). In this annealed regime each cluster sees a fixed rate s_0 but the failure hazard is averaged over the

stationary distribution of s_0 before a failure occurs, so

$$\frac{T_{\text{mem}}}{T_{\text{mem}}^{\text{flat}}} = \frac{h(1)}{\langle h(s_0) \rangle}, \quad \tau_{\text{cluster}} \ll \tau_c \ll T_{\text{mem}}, \quad (2)$$

with $h(s_0) = 1/T_{\text{mem}}(s_0)$ the single-rate hazard. For a power law $h \propto s_0^a$ with $a = 0.6$ (the diffusion sensitivity) this gives $e^{a(1-a)\sigma^2/2} = 1.13$ ($\sigma = 1$) and 1.62 ($\sigma = 2$); using the measured hazard curve of Table I (including the $s_0 = 0.001$ point) and log-log interpolation gives 1.18 and 1.78, against the measured 1.1–1.2 and 2.0–2.1 ($\tau_c = 10^2$ – 10^6). Two independent checks support the annealed assignment: the opposite, quenched limit $\tau_c \gg T_{\text{mem}}$ would

give $\langle 1/h \rangle / (1/h(1)) = e^{a(1+a)\sigma^2/2} = 6.8$ for $\sigma = 2$, excluded by the data at $\tau_c = 10^6$ (1.97 ± 0.26), and a quenched mixture of exponentials would have a coefficient of variation of 2.7, whereas the measured value is 0.94 ± 0.09 (the error being the sampling standard deviation of a sample CV at $n = 100$ under an exponential law). To test the opposite limit directly we ran $\tau_c = 10^{10}$, twelve times the resulting memory time, with ξ drawn from its stationary distribution ($n = 200$): the gain is 7.0 ± 1.1 , against the quenched power-law prediction 6.8, and the coefficient of variation of the lifetimes rises to 1.72. Three further quenched points ($n = 200$ each) test the dependence on σ and on size: at $\sigma = 1$, $L = 9$ the gain is 1.63 ± 0.20 against the predicted $e^{a(1+a)/2} = 1.62$, and at $\sigma = 2$ it is 10.9 ± 2.0 at $L = 7$ and 6.1 ± 0.8 at $L = 11$ against 6.8. With the single exponent $a = 0.6$ the four points give $\chi^2 = 5.1$ (no fitted parameter), the $L = 7$ point lying 2.1 standard errors above the prediction. The long-correlation points therefore have means consistent with the quenched power-law estimate. Its coefficient of variation must be compared with the sampling distribution of a sample CV at $n = 200$, not with the population value: under the fully quenched power-law model the sample CV has median 2.14 with 5–95% range [1.62, 3.58], so the measured 1.72 is compatible with the quenched law (at its 11th–12th percentile) and incompatible with an exponential law (probability below 10^{-5} ; Appendix B5). The measured-hazard version of the quenched average, 10.5 when the archived $s_0 = 0.001$ point is included (13.6 without it), overshoots the observation; the measured hazard is not a single power law (local slopes 0.54–0.82 between successive decades), and the instantaneous-hazard approximation, finite τ_c/T_{mem} , and finite-sample tail fluctuations remain candidate explanations that the present data do not separate. Equation (2) and its quenched counterpart (Appendix B) are conditional explanations, not theorems. The gain at $\sigma = 2$, $\tau_c = 100$ is 1.96 ± 0.24 , 1.75 ± 0.27 , 2.12 ± 0.30 , 1.69 ± 0.21 , 1.36 ± 0.18 for $L = 5, 7, 9, 11, 13$, with a weighted slope $d \ln(\text{gain})/d \ln L = -0.29 \pm 0.18$: a mild decrease with size is neither established nor excluded.

D. A locally damped structured mode: populations by identity, rates by dynamics

For the one-defect pseudomode model the excitation number $Q = \sigma^+ \sigma^- + a^\dagger a$ commutes with the rotating-wave Hamiltonian at any detuning, and the product state $\rho_* \propto e^{-\beta \omega_0 Q}$ is annihilated by the local thermal dissipator; it is therefore an exact stationary state of this generator for every g and κ . At resonance its qubit marginal is the bare Boltzmann state, and the numerical steady state reproduces this to 10^{-7} over $g \in [0.05, 1]$, $\kappa \in [0.03, 10]$, $\beta \in \{2, 4.5\}$ (Fig. 4a); the numerics add only that this fixed point is the attractor. Off resonance the same identity gives a qubit population ratio $e^{-\beta \omega_0}$ rather than $e^{-\beta \Delta}$, which is the well-known failure of a local thermal




FIG. 3. **A bath correlation time adds a factor of two while the bath renews within a memory time and a factor of seven at the longest correlation time sampled, consistent with the annealed and quenched hazard averages.** Memory time relative to the flat bath at $L = 9$, $\beta = 4.5$, versus the correlation time τ_c of a log-normal Ornstein–Uhlenbeck modulation of the diffusion rate with mean one and log-variance σ^2 . Dashed lines: annealed limit of Eq. (2) with $h \propto s_0^{0.6}$; dotted horizontal line: quenched limit $\langle s_0^{-0.6} \rangle$ for $\sigma = 2$; dotted vertical line: the flat-bath memory time. Error bars include the flat-bath denominator; $n = 100$ runs per point (200 at $\tau_c = 10^{10}$).

Lindblad term to thermalize a detuned dressed system; a locally damped mode therefore cannot be used to ask whether a detuned structured bath changes defect populations, and we do not draw conclusions about general quantum baths from it. What the model does determine is the kinetics: the population relaxation rate is the golden-rule value $4g^2/\kappa$ to within 0.5% for $\kappa \geq 50g$, deviates by 4% at $\kappa = 10g$, and saturates at $\kappa/2$ for $g \gtrsim \kappa/3$, when the excitation is shared coherently between qubit and mode. A finite bath memory time thus caps the rates in this model without changing its stationary population.

The equilibrium question can be asked without dynamics through the exact reduced Gibbs state of qubit plus mode. Its bare-defect population p_{def} (Fig. 4b) is not below the Boltzmann value at any sampled parameter: for spectral weight above the gap it follows the second-order shift of the qubit-like level, $p_{\text{def}}/p_B \simeq e^{\beta g^2/(\omega_0 - \Delta)}$ (1.96 measured against 2.25 predicted at $\omega_0 = 6$, $g = 0.6$, $\beta = 4.5$), while for spectral weight below the gap it is strongly enhanced (1.2×10^3 at $\omega_0 = 2$, $g = 0.6$; 2.4×10^5 at $\omega_0 = 1$, $g = 1$), because a thermally populated sub-gap mode admixes the excited qubit into eigenstates of low energy. The qubit-like eigenvalue itself shifts by the exact Jaynes–Cummings amount $\frac{1}{2}[\sqrt{D^2 + 4g^2} - |D|] \text{sgn}(D)$ with $D = \Delta - \omega_0$ (Appendix B7) (0.044, 0.166, 0.414 for $g = 0.3, 0.6, 1.0$ at $D = 2$; the second-order value g^2/D overestimates the last by 20%), which is a renormalization of the spectrum, not of the population. Whether

TABLE II. **A gapped structured mode does not act as a uniform energy rescaling: the mean-force defect energy shifts with the wrong sign below the gap and a shared mode induces defect-defect interactions.** Two defect qubits (gap $\Delta = 4$) coupled with rotating-wave coupling g to one shared pseudomode of frequency ω_0 , $\beta = 4.5$; single-defect mean-force shift $\epsilon_1 - \Delta$, induced interaction U_{12} , and the Jaynes–Cummings shift of the qubit-like level for comparison. With independent modes $U_{12} = 0$ and $\epsilon_1 - \Delta$ changes by less than 0.25.

ω_0	g	$\epsilon_1 - \Delta$	U_{12}	JC level shift
2	0.3	-1.22	-0.14	+0.04
2	0.6	-1.68	-0.11	+0.17
2	1.0	-2.23	-0.03	+0.41
6	0.3	-0.04	+0.00	-0.04
6	0.6	-0.18	+0.05	-0.17
6	1.0	-0.54	+0.22	-0.41

the defects relevant to the code’s kinetics are the bare or the dressed ones, and whether the admixture is removed at readout, are questions this one-defect model cannot answer. A two-defect extension tests the minimal condition for the environment to act as a uniform energy rescaling, namely that the mean-force energy of a defect shifts additively with no induced interaction. With two defect qubits coupled to one shared mode (or to two independent modes), we compute the sector free energies $F_{ij} = -\beta^{-1} \ln \text{Tr}[\Pi_{ij} e^{-\beta H}]$ and from them the single-defect shift $\epsilon_1 - \Delta = F_{10} - F_{00} - \Delta$ and the induced interaction $U_{12} = F_{11} - F_{10} - F_{01} + F_{00}$ (Table II). For sub-gap spectral weight the shift is negative and large, -1.2 to -2.2 defect energies at $\omega_0 = 2$, opposite in sign to the Jaynes–Cummings level shift, so in mean-force terms a sub-gap bath makes defects cheaper, not dearer; for spectral weight above the gap the shift follows the level shift in sign and roughly in magnitude (-0.18 against -0.17 at $\omega_0 = 6$, $g = 0.6$). A shared mode induces an interaction U_{12} between -0.14 and $+0.22$ for $g \geq 0.3$, attractive below and repulsive above the gap, while independent modes give $U_{12} = 0$ to numerical precision. The sampled models therefore do not realize the positive uniform rescaling assumed in the benchmark of Sec. III E: with a shared mode the shift is non-additive, and with independent modes it is additive but negative below the gap. These are statements about occupation-sector free energies of gap-4 toy defects, not about the equilibrium of the code, whose coupling operators and syndrome constraints are absent here. A related caveat is that a counter-rotating coupling with the same local damping converts the ground-state dressing $\epsilon = g/(\Delta + \omega_0)$ into real excitations at rate $\kappa\epsilon^2$, giving a spurious population $\kappa\epsilon^2/\Gamma$ of 10^5 times the Boltzmann value at $\beta = 4.5$; any quantitative study of populations in such models must use a thermalizing global generator.

E. Uniform energy rescaling as a cooling benchmark

If an environment renormalized every code energy by a common factor $c = 1 + \delta/\Delta$ with unchanged kinetic prefactors, the Davies rates would be those of the flat bath at $\beta' = c\beta$. The two-defect calculation above shows that the gapped pseudomode does not provide such a rescaling (the mean-force shift has the wrong sign below the gap and a shared mode induces interactions), so the following is a benchmark of what a uniform rescaling would buy if some environment provided it, not a demonstration that one does. Figure 5 and Table III show the flat-bath memory time at $\beta' = 4.70, 4.875, 5.06$, i.e. $\delta = 0.18, 0.33, 0.50$ (in defect-energy units; 12.5% of the four-defect gap for the largest). On the common window the exponent rises from 2.73 ± 0.17 to 4.34 ± 0.25 , within 0.6σ of the Bravyi–Haah form $2.93\beta' - 10.52$ at every β' , and $T_{\text{mem}}(13)/T_{\text{mem}}(5)$ grows from 9.9 to 87. The size diagnostic for $L = 11 \rightarrow 13$ rises from 1.12 ± 0.14 to 2.80 ± 0.35 . We quote the active-memory value $\Lambda_{\text{active}} = 2.14$ of Ref. [1] in Fig. 5b only as a numerical reference: the two quantities are not commensurable, since Λ_{passive} is a finite-difference of a polynomial that decays to one as $L \rightarrow L^*$ ($L^* = 21.7$ at $\beta' = 5.06$, four size steps beyond our data), a step $L \rightarrow L + 2$ costs 1732 physical qubits at $L = 11$ against 64 for a surface-code step $d = 7 \rightarrow 9$, and the two memories are evaluated with different noise models, clocks, and success criteria. No conclusion about passive versus active performance follows from Table III.

IV. DISCUSSION

Three statements summarize the benchmark, each with its scope.

First, within a reversible bath the equilibrium abundance of defect clusters is set by the code and the temperature, and over the tested range the spectral weight at the pair-creation energy, the quantity a phononic band gap would most naturally suppress, does not change the memory time. The scope is the escape-limited regime $s_4 \gg s_4^*$; below s_4^* the two additional points show the memory time rising as the first-error time, consistent with an additive form that we use as an interpolation, not as a derived law. For self-correcting memories with parity-check redundancies the syndrome sector equilibrates rapidly within a logical sector under a quasi-local Gibbs sampler [46]; we note that rapid equilibration does not by itself fix first-passage times, and that the cubic code is outside the class treated there. Conversely, for expander-based low-density parity-check codes without redundant checks the Metropolis mixing time diverges below a critical temperature while the free energy remains analytic [47]: there slow mixing is itself the protection and no thermodynamic transition is needed. The cubic code, whose lifetime is finite at any $T > 0$ because




FIG. 4. **A locally damped resonant mode leaves the defect population at its Boltzmann value by an exact stationary-state identity; the reduced Gibbs population of a bare defect is enhanced, not suppressed, by sub-gap spectral weight.** (a) Steady-state excited population of a defect qubit (gap $\Delta = 4$) with rotating-wave coupling g to a resonant mode damped at rate κ , relative to the Boltzmann value, versus $\kappa = 1/\tau_c$. (b) Bare-defect population $\text{Tr}[\Pi_e \rho_{\text{Gibbs}}]$ of the exact reduced Gibbs state of qubit plus mode, relative to Boltzmann, versus the mode frequency ω_0 ; the resonant point is included. Dashed and solid curves: $\beta = 2$ and 4.5.

TABLE III. **A uniform rescaling of the code energies moves the size exponent along the Bravyi–Haah line.** Flat-bath memory time at the effective inverse temperature $\beta' = \beta(1 + \delta/\Delta)$, $\beta = 4.5$, $n = 100$; exponent on $L = 5$ –11 with scaled error, χ^2/dof and window systematic; the Bravyi–Haah value $2.93\beta' - 10.52$ (BH); and the size diagnostic of Eq. (1).

δ	β'	exponent	$\chi^2/\text{dof}; \Delta p_{\text{win}}$	BH	$\Lambda_{\text{passive}}(9 \rightarrow 11)$	$\Lambda_{\text{passive}}(11 \rightarrow 13)$
0	4.50	2.73(17)	1.2; 0.37	2.67	1.69(21)	1.12(14)
0.18	4.70	3.33(17)	0.7; 0.03	3.25	2.29(33)	1.64(23)
0.33	4.875	3.88(22)	1.9; 0.08	3.76	2.77(35)	1.77(24)
0.50	5.06	4.34(25)	2.7; 0.45	4.31	2.89(38)	2.80(35)




FIG. 5. **A uniform energy rescaling of 12% raises the size exponent from 2.7 to 4.3, as the Bravyi–Haah form predicts.** (a) Memory time versus L at the effective inverse temperatures $\beta' = \beta(1 + \delta/\Delta)$; dotted vertical lines mark $L^*(\beta')$. (b) Size diagnostic $\Lambda_{\text{passive}} = T_{\text{mem}}(L+2)/T_{\text{mem}}(L)$ versus L ; the dashed line is the surface-code value $\Lambda_{\text{active}} = 2.14$ of Ref. [1], shown as a numerical reference only (see text). Error bars are standard errors over $n = 100$ runs.

of entropic effects, belongs to neither class, and the first-failure times of Figs. 1 and 2 are the finite-size kinetics of that intermediate case.

Second, the kinetic lever is the zero-energy rate, and on the present data it acts as a prefactor. That a vanishing zero-frequency thermal rate freezes anyon diffusion was established for the toric code in Ref. [21], where, with engineered long-range repulsion, it also improved the size scaling; in the cubic code the same mechanism buys a factor of 78–121 for $L = 7$ –13 (40 at $L = 5$) over three decades of s_0 ; for $s_0 \leq 0.01$ the size dependence is not a single power law and is weaker than in the flat bath beyond $L = 7$, so the gain is a prefactor and not a steeper size law (Fig. 2a). This is consistent with the mechanism that caps translation-invariant fractal codes, the multiplicity of low-energy paths along which a defect cluster can spread [19], which has been formalized for stabilizer Hamiltonians as a free-energy barrier that bounds the thermalization time [48]: a filter on the zero-energy rate slows every such path by the same factor without changing their number or the barrier, and the bound of Ref. [48], which is asymptotic and holds the kinetic rates between fixed constants, neither predicts nor excludes the finite-size turnover of Fig. 1. A bath correlation time adds a factor of two in the annealed regime and about seven at the longest correlation time sampled, consistent with Eq. (2) and its quenched counterpart with the same exponent, the lifetime distribution at that point being compatible with the quenched law at the level of its sampling distribution; we are not aware of a previous treatment of a self-correcting memory in a temporally correlated bath (the closest works concern active correction in correlated environments [26, 49] and, most recently, continuous measurement-based correction of small codes, where a bath memory improves the fidelity through a Zeno regime [50]; a passive memory has no measurement, and the gain in Fig. 3 has the different, kinetic origin of Eq. (2)).

Third, a uniform rescaling of the code energies moves the exponent along the Bravyi–Haah line, so any environment that produced such a rescaling would act as cooling. Bath-induced anyon penalties are known: massless bosons give a penalty growing with L [24, 25], gapped bosons a finite one [14], and a finite-size memory enhanced by an inverse Boltzmann factor is a resource even when it fails the criteria for self-correction [14]. Our pseudomode models show a finite spectral shift of the qubit-like level but also that the occupation-sector free energy of a bare toy defect decreases, and that a shared mode induces defect–defect terms, for sub-gap spectral weight; the sampled models therefore do not realize the positive rescaling of Table III, and whether a gapped structured bath cools or heats the code itself is not determined by them. Non-equilibrium routes, such as driven-dissipative ancillas that trap defects [28], a continuously running cellular-automaton decoder [30], or dissipative phases that do not rely on detailed balance [32], and disorder-induced localization [51, 52], lie outside the re-

versible class studied here.

The limitations delimit these statements. All lifetimes are first-failure times of one RG decoder with a generic neutrality test; efficient decoders for fracton codes are largely unknown, and even the code-capacity threshold of the cubic code is at present only conjectured, from an entropy duality, to lie near 0.11 [53], so the failure counts reported here characterize this decoder, not the code. The fitted series reach $L = 13$, below L^* at every temperature ($L^* = 14.0$ at $\beta = 4.5$), so exponents are finite-window quantities with a window systematic of 0.03–0.45; only the flat bath and the $s_0 = 0.01$ series extend beyond the optimum, and only at $\beta = 4.5$; for the strongest diffusion filter no single power law describes the data. The filters are phenomenological reversible rate multipliers in the X sector; the correspondence to a microscopic spectral density is stated only at the level of the thermal rate, and simultaneous-Pauli or spatially correlated noise [26] is not modeled. The Ornstein–Uhlenbeck modulation is an external, spatially uniform rate process with a fixed time step $\tau_c/20$ whose refinement has not been tested. The pseudomode model has one defect, one mode and a local damping whose bare qubit marginal is thermal only at resonance; it contains neither defect–defect interactions nor a readout channel, and the two-defect extension shows that its spectral shift does not translate into a uniform rescaling of the code energies. Finally all filtered series share random seeds with the flat bath and per-run lifetimes are not archived, so the covariance between series is unknown and cross-series differences are reported with quadrature-combined errors whose direction of bias is not established. The computations that would remove these limitations are: a many-defect mean-force calculation with the code’s coupling operators, extending the two-defect test of Table II; sizes $L = 15$ –17 at the highest β' to reduce the window systematic; and a comparison of failure syndromes against a stronger recovery to separate decoder failure from loss of logical information.

V. CONCLUSION

We have measured, for Haah’s cubic code, how the memory time of a passive three-dimensional memory responds to each of three phenomenological environmental controls, providing to our knowledge the first channel-resolved benchmark of this kind for a three-dimensional code. The memory time responds to the zero-energy rate with logarithmic sensitivity -0.6 over three decades and to a bath correlation time by a factor of two (annealed) to seven (quenched) consistent with a hazard average, while the pair-creation rate is invisible on the plateau $s_4 \in [0.01, 10]$ and takes over below the creation-limited crossover, and a uniform energy rescaling of 12% raises the fitted size exponent from 2.7 to 4.3 along the Bravyi–Haah line. Compared with the interacting toric code of Ref. [21], where the same zero-frequency mechanism improves the scaling, the cubic code gains a factor of 13–45

at $s_0 = 0.01$ over $L = 5\text{--}19$ (and $40\text{--}121$ at $s_0 = 0.001$ over $L = 5\text{--}13$) with no change of exponent beyond the tolerance for the series a power law describes and no steeper growth for those it does not; compared with a flat Markovian bath, a locally damped resonant mode changes rates but, by an exact identity, not populations. For the broader effort to build memories that do not need a decoder, these numbers locate the environment among the design variables: at fixed code and temperature the zero-frequency spectral weight is the lever worth engineering, temporal correlations are worth a factor of two to seven at the sampled amplitudes, and among the controls tested, only a uniform rescaling of the code energies moved the fitted size exponent of a series that a power law describes (the strongest diffusion filter, which no power law describes, grows more slowly than the flat bath on $L = 7\text{--}13$), and no bath in the reversible class studied here, including the gapped structured mode of the two-defect model, has been shown to provide one. On the present data a bath correlation time is worth a factor of two (annealed) to six to eleven (quenched, $L = 7\text{--}11$), and a suppressed zero-frequency thermal rate, as in a super-Ohmic phonon bath, is worth up to two orders of magnitude, and at $s_0 = 0.01$ the gain survives and grows beyond the flat-bath optimum, which the flat bath crosses between $L = 13$ and 17 at $\beta = 4.5$ while the filtered bath still grows to $L = 17$; whether a bath correlation time also does is open. A many-defect mean-force calculation with the code's coupling operators, and the location of the shifted optimum under a suppressed diffusion rate, are the immediate next steps. Beyond the cubic code, the same channel decomposition applies to any code with a classical Davies reduction, including the layer codes [16, 17] and the translation-invariant families obtained by compactifying fracton models [54], whose thermal stability is so far conjectural.

ACKNOWLEDGMENTS

Numerical simulations were performed with custom Python and numba code.

DATA AVAILABILITY

The code construction, decoder, simulation scripts with seed and decoder-cap arguments, all result files with a provenance manifest, the figure scripts, and a test suite that regenerates the derived numbers, the figures, and two archived Monte Carlo points bit for bit are available at <https://github.com/WakauraH/haah-cubic-bath-benchmark>.

Appendix A: Validation of the simulation

Exact master equation at $L = 2$. For $N = 16$ the classical Davies chain on all 2^{16} error configurations can be integrated exactly by uniformization, $P(t) = \sum_n \text{Pois}(n; \Lambda t) (I + W/\Lambda)^n P(0)$ with $\Lambda = N\gamma_0$. In this check the Hamiltonian is written as $-\Delta_s \sum_c S_c$ with eigenvalues $\pm\Delta_s$ per stabilizer, so that a single flip, which violates four stabilizers, costs $8\Delta_s$, and Glauber rates $\gamma_0/(1 + e^{\beta\Delta_s})$ are used. With a minimum-weight lookup decoder the Gillespie simulation reproduces the exact success probability with maximum deviation $0.014\text{--}0.040$ over $\beta = 0.5\text{--}1.5$, and the memory time scales as $e^{7.8\beta\Delta_s}$, consistent with the $8\Delta_s$ single-flip barrier; the memory times themselves agree exactly on the time grid at $\beta = 0.5, 1.0$ and 1.25 and differ by 23% at $\beta = 0.75$ and 1.5 , where the threshold crossing falls between grid points; this check uses a minimum-weight lookup decoder and the first grid time at which the success probability drops below $1 - 1/e$, not the trial-decoding protocol of the main text.

Decoder. At $L = 5$ every single-qubit error is corrected; at $L = 9$ all of 100 random weight-two and 100 weight-three errors are corrected with $L_{\text{tqo}} = L - 2$ (96% of the weight-three errors with $L_{\text{tqo}} = L/2$, where three nearby errors can merge into a cluster whose box exceeds the cap; both settings are exposed as a script argument). With $L_{\text{tqo}} = L/2$ the memory times at $L = 7$ and 9 are identical to those with $L - 2$ for the same random seeds and differ only from $L = 11$ upward.

Time resolution and statistics. The trial-decoding period $T_{\text{ec}} = e^{4\beta}/100$ must be small compared with T_{mem} ; at $\beta = 4.0$ the archived ratio is only $15\text{--}21$ ($L = 5, 7$; $n = 10$), too coarse for the per-size comparisons made here, which is why the many-body scans are reported for $\beta \geq 4.5$, where $T_{\text{mem}}/T_{\text{ec}}$ ranges from 34 ($L = 5$) to 8.8×10^3 ($L = 13$, $\beta' = 5.06$). An independent-seed rerun of the flat bath ($n = 60$, seeds offset by 5×10^5 , archived as `tmem_beta4.5_reviewA_flat.json`) gave $8.0(11) \times 10^7$ at $L = 7$ and $1.20(14) \times 10^8$ at $L = 9$, against the tabulated $7.1(8)$ and $1.20(11) \times 10^8$. Two result files were overwritten by later runs and regenerated with the same seeds, as recorded in the manifest: the $L = 9$ entry of the $s_0 = 0.01$ series reproduced its logged value exactly, and the $\sigma = 2$ Ornstein–Uhlenbeck records at $\tau_c = 1$ and 10^2 did likewise, while $\tau_c = 10^6$ moved from 2.44 to 2.36×10^8 because the initial value $\xi(0)$ is now drawn from the seeded generator; the earlier records (and the $\tau_c = 10^{10}$ run) drew $\xi(0)$ from an unseeded generator and are statistically valid but not bit-reproducible. The oscillator cutoff sweeps of both pseudomode models and the $\sigma \rightarrow 0$ self-test of the modulated kernel are archived with the results.

Appendix B: Derivations used in the text

Each derivation lists its inputs (starting equations with their primary sources) and then numbered steps, so that every result used in the main text can be followed without external material.

1. Classical reduction of the Davies dynamics

Inputs. (I1) The Davies generator [41, 55], in the form used in Ref. [12], Eqs. (1)–(3): $\dot{\rho} = -i[H, \rho] + \mathcal{L}(\rho)$ with $\mathcal{L}(\rho) = \sum_{\alpha, \omega} h(\alpha, \omega) (A_{\alpha, \omega} \rho A_{\alpha, \omega}^\dagger - \frac{1}{2} \{A_{\alpha, \omega}^\dagger A_{\alpha, \omega}, \rho\})$, where $A_{\alpha, \omega}$ is the component of the coupling operator A_α that lowers the energy by ω and the rates obey detailed balance $h(\alpha, -\omega) = e^{-\beta\omega} h(\alpha, \omega)$. (I2) The code Hamiltonian $H = -J \sum_c (G_c^X + G_c^Z)$ with mutually commuting stabilizers, $J = 1/2$ [11, 12]. (I3) Coupling operators $A_\alpha = X_q$, one per qubit.

Steps. (1) Since all G_c^Z commute, the joint eigenbasis of $\{G_c^Z\}$ diagonalizes the Z -stabilizer part of H ; an X -error configuration $e \in \{0, 1\}^N$ applied to a code state gives the eigenstate $|e\rangle$ with energy $E(e) = \sum_c [G_c^Z(e) = -1]$ in defect units. Two frames that differ by an X -stabilizer give the same state, so the chain on frames e is a lift of the physical dynamics; the lift is harmless because the energy, the syndrome, and the logical class of e depend only on e modulo the stabilizer group, and readout uses only those. (2) X_q maps $|e\rangle$ to $|e \oplus q\rangle$ and flips exactly the four stabilizers containing q , so $E(e \oplus q) - E(e) = 4 - 2n_{\text{viol}}(q)$; hence the spectral component of X_q with frequency $\omega = -\Delta E$ is X_q restricted to the configurations with the corresponding $n_{\text{viol}}(q)$. (3) For ρ diagonal in the $|e\rangle$ basis, $[H, \rho] = 0$ and each term of $\mathcal{L}$ maps $|e\rangle\langle e|$ to a combination of $|e\rangle\langle e|$ and $|e \oplus q\rangle\langle e \oplus q|$; the generator therefore preserves diagonality and its restriction to populations $P(e)$ is a classical master equation with transition rate $W(e \rightarrow e \oplus q) = h(q, -\Delta E)$. (4) Choosing the Metropolis form $h(q, -\Delta E) = \gamma_0 \min(1, e^{-\beta\Delta E})$ [42, 43] satisfies (I1)'s detailed balance, $W(e \rightarrow e')/W(e' \rightarrow e) = e^{-\beta[E(e') - E(e)]}$. (5) The rejection-free Bortz–Kalos–Lebowitz algorithm [44] samples this chain exactly: with total rate $R = \sum_q W(e \rightarrow e \oplus q)$ the waiting time is exponential with mean $1/R$ and the flipped qubit is drawn with probability W/R . Ergodic decompositions of Davies generators for CSS codes, which justify treating the X sector alone, are given in Ref. [13].

2. Invariance of the Gibbs distribution under channel filters

Inputs. (I1) The classical chain of Appendix B 1 with detailed balance $\pi(e)W(e \rightarrow e') = \pi(e')W(e' \rightarrow e)$, $\pi(e) \propto e^{-\beta E(e)}$. (I2) The filter definition: $W \rightarrow s_{|\Delta E|} W$ for both directions of every transition with energy change $\pm \Delta E$.

Steps. (1) Multiplying both sides of the detailed-balance identity by the same $s_{|\Delta E|} > 0$ leaves it valid with the same π . (2) Hence π is stationary for the filtered chain and, the chain being irreducible for $s > 0$, unique. (3) This fixes stationary occupation probabilities only; first-passage times (the memory time) are not determined by π , as the bound of Appendix B 3 shows.

3. First-error bound and creation-limited crossover

Inputs. (I1) From the code state every qubit has $n_{\text{viol}} = 0$, so every flip has $\Delta E = +4$ and rate $s_4 \gamma_0 e^{-4\beta}$ (Appendix B 1, step 2 and 4). (I2) The readout protocol can count a failure only if at least one error is present, since a zero syndrome with $e = 0$ decodes trivially.

Steps. (1) The first flip occurs after a random time τ_{first} , exponential with mean $T_{\text{first}} = 1/(Ns_4\gamma_0 e^{-4\beta})$, $N = 2L^3$. (2) By (I2) no failure precedes the first flip, so the memory time dominates τ_{first} in distribution and $\mathbb{E} T_{\text{mem}} \geq T_{\text{first}}$. (3) Defining s_4^* by $T_{\text{first}}(s_4^*) = T_{\text{mem}}^{\text{flat}}$ gives $s_4^* = 1/(Ne^{-4\beta}\gamma_0 T_{\text{mem}}^{\text{flat}}) = 3.8 \times 10^{-4}$ at $L = 9$, $\beta = 4.5$. (4) The additive form $T_{\text{mem}} \simeq T_{\text{mem}}^{\text{flat}} + T_{\text{first}}$ used in Sec. III B is an interpolation motivated by the two limits, not a consequence of steps 1–3.

4. Uniform energy rescaling as a temperature change

Inputs. (I1) The Metropolis rates $\gamma(\Delta E) = \gamma_0 \min(1, e^{-\beta\Delta E})$ of Appendix B 1. (I2) The hypothesis $H \rightarrow cH + \text{const}$ with $c = 1 + \delta/\Delta$ and unchanged jump operators and prefactor γ_0 .

Steps. (1) Under (I2) every $\Delta E \rightarrow c\Delta E$, so $\beta\Delta E \rightarrow c\beta\Delta E = \beta'\Delta E$ with $\beta' = c\beta$. (2) The chain at (H, β') with the original energies is therefore identical to the chain at (cH, β) , transition by transition. (3) Step 1 requires that all channels, $|\Delta E| = 2$ and 4, are rescaled by the same c and that no configuration-dependent terms are added; Sec. III D and Table II test and reject this hypothesis for the pseudomode models.

5. Hazard averages for a modulated diffusion rate

Inputs. (I1) At fixed diffusion rate s_0 the lifetime is taken to be exponential with hazard $h(s_0) = 1/T_{\text{mem}}(s_0)$ (empirical, Sec. II C: coefficient of variation 0.97 ± 0.09 over the 79 fixed-rate points). (I2) The modulation $s_0(t) = \exp[\xi(t) - \sigma^2/2]$ with ξ a stationary Ornstein–Uhlenbeck process [45], $\xi(0) \sim \mathcal{N}(0, \sigma^2)$, correlation time τ_c ; hence $\langle s_0 \rangle = 1$. (I3) The instantaneous-hazard approximation: conditional on a rate history the survival probability is $\exp[-\int_0^t h(s_0(u)) du]$, the survival function of a process with time-dependent hazard [56], and the unconditional survival is its average over histories, $S(t) = \mathbb{E} \exp[-\int_0^t h(s_0(u)) du]$. This treats the failure

hazard as a function of the instantaneous rate alone, which is an approximation for a process whose failures are built up over a cluster lifetime.

Steps. (1) *Annealed limit* $\tau_{\text{cluster}} \ll \tau_c \ll T_{\text{mem}}$: by ergodicity of the stationary process the time average $t^{-1} \int_0^t h(s_0(u)) du$ converges to $\langle h \rangle$ for $t \gg \tau_c$ (homogenization of the hazard), so $S(t) \simeq e^{-\langle h \rangle t}$ on the scale of T_{mem} and $\mathbb{E}T = 1/\langle h(s_0) \rangle$, which is Eq. (2). (2) *Quenched limit* $\tau_c \gg T_{\text{mem}}$: s_0 is constant during a run and drawn from its stationary law, $S(t) = \langle e^{-h(s_0)t} \rangle$ and $\mathbb{E}T = \int_0^\infty S dt = \langle 1/h(s_0) \rangle$. (3) *Log-normal moments*: for $s_0 = e^{\xi - \sigma^2/2}$ with $\xi \sim \mathcal{N}(0, \sigma^2)$, $\langle e^{a\xi} \rangle = e^{a^2\sigma^2/2}$ gives $\langle s_0^a \rangle = e^{a(a-1)\sigma^2/2}$. With $h \propto s_0^a$: annealed gain $1/\langle s_0^a \rangle = e^{a(1-a)\sigma^2/2}$, quenched gain $\langle s_0^{-a} \rangle = e^{a(1+a)\sigma^2/2}$; for $a = 0.6$, $\sigma = 2$ these are 1.62 and 6.82. (4) *Quenched coefficient of variation*: $\mathbb{E}T^2 = 2\langle h^{-2} \rangle$, so $\text{CV}^2 = 2\langle s_0^{-2a} \rangle / \langle s_0^{-a} \rangle^2 - 1 = 2e^{a^2\sigma^2} - 1$, i.e. $\text{CV} = 2.73$ for $a = 0.6$, $\sigma = 2$; the annealed limit has $\text{CV} = 1$. (5) *Measured-hazard version*: $h(s_0)$ is taken at $s_0 \in \{0.001, 0.01, 0.1, 1, 10\}$ ($L = 9$), interpolated linearly in $\log h$ versus $\log s_0$ and extrapolated with the end slopes; the averages are Monte Carlo integrals over 2×10^6 samples of ξ (script `protocol.f.py`). Including the archived $s_0 = 0.001$ point lowers the quenched measured-hazard value from 13.6 to 10.5, so the extrapolation accounts for part but not all of the overshoot discussed in Sec. III C. (6) *Sampling distribution of the CV*: for the fully quenched model with $a = 0.6$, $\sigma = 2$ and $n = 200$, 4×10^5 replicas of T_i exponential with mean $\propto s_{0,i}^{-a}$ (the constant cancels in the CV), with independent log-normal $s_{0,i}$, give a sample CV with median 2.14 and 5–95% range [1.62, 3.58]; for an exponential law the sample CV at $n = 200$ never reaches 1.72 in the same number of replicas, and its sampling standard deviation at $n = 100$ is 0.095, quoted as 0.09 in Sec. III C (script `protocol.f.py`).

6. Stationary state of the locally damped rotating-wave model

Inputs. (I1) $H = \frac{\Delta}{2}\sigma^z + \omega_0 a^\dagger a + g(\sigma^+ a + \sigma^- a^\dagger)$ [57]. (I2) The local thermal dissipator of a damped oscillator, $\mathcal{D}(\rho) = \kappa(n+1)\mathcal{D}[a]\rho + \kappa n\mathcal{D}[a^\dagger]\rho$ with $\mathcal{D}[c]\rho = c\rho c^\dagger - \frac{1}{2}\{c^\dagger c, \rho\}$ and $n = (e^{\beta\omega_0} - 1)^{-1}$, whose unique stationary state for the oscillator alone is the thermal state $\rho_{\text{th}} \propto e^{-\beta\omega_0 a^\dagger a}$ [55]. (I3) $Q = \sigma^+ \sigma^- + a^\dagger a$.

Steps. (1) Each term of (I1) conserves Q , so $[H, Q] = 0$ at any detuning. (2) $\rho_* = e^{-\beta\omega_0 Q}/Z = (e^{-\beta\omega_0 \sigma^+ \sigma^-}/Z_q) \otimes \rho_{\text{th}}$ with $Z_q = 1 + e^{-\beta\omega_0}$ factorizes because $\sigma^+ \sigma^-$ and $a^\dagger a$ commute. (3) $[H, \rho_*] = 0$ by step 1, and $\mathcal{D}$ acts only on the oscillator factor, which it annihilates by (I2); hence $\mathcal{L}(\rho_*) = 0$ exactly for every g , $\kappa > 0$, $\omega_0 > 0$ and $0 < \beta < \infty$. (4) The qubit marginal has excited-to-ground ratio $e^{-\beta\omega_0}$, which for finite positive β equals the Boltzmann ratio $e^{-\beta\Delta}$ if and only if $\omega_0 = \Delta$. (5) Uniqueness (that ρ_* is the attractor) is not

implied by steps 1–4; it was checked numerically over the sampled grid ($\beta \in \{2, 4.5\}$, $g \in [0.05, 1]$, $\kappa \in [0.03, 10]$, cutoff 12): each Liouvillian has exactly one eigenvalue with $|\lambda| < 10^{-10}$ and a spectral gap of at least $5.0e - 04$ (script `pseudomode.uniqueness.py`).

7. Jaynes–Cummings level shift

Inputs. (I1) H of Appendix B 6 restricted to the one-excitation subspace $\{|e, 0\rangle, |g, 1\rangle\}$, where it is the 2×2 matrix with diagonal $(\Delta/2, -\Delta/2 + \omega_0)$ and off-diagonal g [57].

Steps. (1) With $D = \Delta - \omega_0$ the two eigenvalues are $\omega_0/2 \pm \frac{1}{2}\sqrt{D^2 + 4g^2}$. (2) For $D \neq 0$ the eigenvalue that continues $|e, 0\rangle$ as $g \rightarrow 0$ is shifted from $\Delta/2$ by $\frac{1}{2}[\sqrt{D^2 + 4g^2} - |D|] \text{sgn}(D)$, i.e. upward for $D > 0$ (mode below the gap) and downward for $D < 0$. At $D = 0$ no branch continues $|e, 0\rangle$ and the levels split symmetrically by $\pm g$, which is why the resonant point is excluded from the shift comparison. (3) Expanding for $g \ll |D|$ gives g^2/D , the second-order value quoted in Sec. III D; at $g = 1$, $D = 2$ the exact shift is 0.414 against 0.500.

8. Two-defect sector free energies

Inputs. (I1) Two qubits of gap Δ coupled to one shared mode, $H = \frac{\Delta}{2}(\sigma_1^z + \sigma_2^z) + \omega_0 a^\dagger a + g \sum_i (\sigma_i^+ a + \sigma_i^- a^\dagger)$, or to two independent modes with the corresponding two-mode Hamiltonian. (I2) The occupation projectors $\Pi_{ij} = \Pi_i^{(1)} \otimes \Pi_j^{(2)} \otimes I_{\text{modes}}$, $i, j \in \{0, 1\}$.

Steps. (1) Define sector free energies $F_{ij} = -\beta^{-1} \ln \text{Tr}[\Pi_{ij} e^{-\beta H}]$; these are the free energies of the bare occupation sectors, not the spectrum of the Hamiltonian of mean force $-\beta^{-1} \ln \text{Tr}_{\text{modes}} e^{-\beta H}$ (the potential of mean force of Ref. [58]; see Ref. [59] for its use in strong-coupling thermodynamics), which may carry off-diagonal terms in the one-excitation sector. (2) The single-defect cost is $\epsilon_1 = F_{10} - F_{00}$ and the induced two-body term $U_{12} = F_{11} - F_{10} - F_{01} + F_{00}$. (3) For independent modes H and $e^{-\beta H}$ factorize over the two qubit–mode pairs, so $\text{Tr}[\Pi_{ij} e^{-\beta H}]$ factorizes and $U_{12} = 0$ identically. (4) A uniform rescaling of the code energies would require $\epsilon_1 = \epsilon_2 = \Delta(1 + \delta/\Delta)$ with $\delta > 0$ and $U_{12} = 0$; Table II tests both.

9. Fit protocol

Inputs. (I1) Sample means T_i and standard errors σ_i at sizes L_i . (I2) The model $\ln T = c + p \ln L$.

Steps. (1) To first order in σ_i/T_i (delta method) the standard error of $\ln T_i$ is σ_i/T_i ; the observations at different L are treated as uncorrelated; weights are $w_i = (T_i/\sigma_i)^2$. (2) Weighted least squares gives (c, p)

and the covariance $(X^\top W X)^{-1}$; $\chi^2 = \sum_i w_i (\ln T_i - c - p \ln L_i)^2$ with dof = $n - 2$. (3) Following the Particle Data Group convention [60], the error of p is scaled by $\sqrt{\max(1, \chi^2/\text{dof})}$. (4) The window systematic is $|p(5-11) - p(7-13)|/2$. (5) The declared tolerance $|\Delta p| < 0.5$ is a descriptive criterion on point estimates; it is not an equivalence test, and covariance between series sharing seeds is unknown (Sec. II C). (6) $\Lambda_{\text{passive}} = T_{i+1}/T_i$ with error $\Lambda_{\text{passive}} \sqrt{(\sigma_{i+1}/T_{i+1})^2 + (\sigma_i/T_i)^2}$, again assuming zero covariance.

10. Exact solution at $L = 2$ by uniformization

Inputs. (I1) The classical generator W constructed as in Appendix B 1 on all 2^{16} configurations, but with the Glauber rates $\gamma_0/(1 + e^{\beta \Delta E})$ and the energy convention $\pm \Delta_s$ per stabilizer of Appendix A. (I2) The uniformization identity [61]: for any $\Lambda \geq \max_e \sum_{e'} W(e \rightarrow e')$, $e^{Wt} = \sum_{n \geq 0} \text{Pois}(n; \Lambda t) (I + W/\Lambda)^n$, where $I + W/\Lambda$ is a stochastic matrix.

Steps. (1) With $\Lambda = N\gamma_0$ (each of the N qubits con-

tributes at most γ_0) the series has non-negative terms; it is summed past the Poisson mode until every remaining weight is below 10^{-14} , and since the weights decrease monotonically beyond the mode the omitted mass is bounded by 10^{-14} times the number of omitted terms, which is negligible at the precision reported. (2) The success probability at t is the sum of $P(e, t)$ over configurations decoded correctly by the minimum-weight lookup decoder.

Appendix C: Ornstein–Uhlenbeck kernel

Between events ξ evolves by the exact Ornstein–Uhlenbeck transition $\xi(t + \Delta t) = \xi e^{-\Delta t/\tau_c} + \sigma \sqrt{1 - e^{-2\Delta t/\tau_c}} \eta$ with η standard normal [62]. When no qubit has $n_{\text{viol}} = 2$ the diffusion rate is irrelevant and ξ is advanced exactly over the event gap; otherwise the total rate is held fixed over $\min(\Delta t_{\text{event}}, \tau_c/20)$, which is exact for piecewise-constant rates. The limit $\sigma \rightarrow 0$ reproduces the flat-bath memory time at $L = 7$ within errors.

-
- [1] Google Quantum AI and Collaborators, Quantum error correction below the surface code threshold, *Nature* **638**, 920 (2025).
 - [2] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature* **627**, 778 (2024).
 - [3] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, *Nature* **549**, 172 (2017).
 - [4] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).
 - [5] D. Bacon, Operator quantum error-correcting subsystems for self-correcting quantum memories, *Phys. Rev. A* **73**, 012340 (2006).
 - [6] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [7] B. J. Brown, D. Loss, J. K. Pachos, C. N. Self, and J. R. Wootton, Quantum memories at finite temperature, *Rev. Mod. Phys.* **88**, 045005 (2016).
 - [8] S. Bravyi and B. Terhal, A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes, *New J. Phys.* **11**, 043029 (2009).
 - [9] S. Chesi, D. Loss, S. Bravyi, and B. M. Terhal, Thermodynamic stability criteria for a quantum memory based on stabilizer and subsystem codes, *New J. Phys.* **12**, 025013 (2010).
 - [10] R. Alicki, M. Fannes, and M. Horodecki, On thermalization in kitaev’s 2d model, *J. Phys. A* **42**, 065303 (2009).
 - [11] J. Haah, Local stabilizer codes in three dimensions without string logical operators, *Phys. Rev. A* **83**, 042330 (2011).
 - [12] S. Bravyi and J. Haah, Quantum self-correction in the 3d cubic code model, *Phys. Rev. Lett.* **111**, 200501 (2013).
 - [13] K. Siva and B. Yoshida, Topological order and memory time in marginally-self-correcting quantum memory, *Phys. Rev. B* **93**, 235159 (2016).
 - [14] O. Landon-Cardinal, B. Yoshida, D. Poulin, and J. Preskill, Perturbative instability of quantum memory based on effective long-range interactions, *Phys. Rev. A* **91**, 032303 (2015).
 - [15] S. Balasubramanian, M. Davydova, and T.-C. Lin, A passive self-correcting quantum memory in three dimensions, *arXiv:2605.10943* (2026).
 - [16] D. J. Williamson, Partial self-correction in layer codes, *arXiv:2510.09218* (2025).
 - [17] S. Gu, L. Caha, S. H. Choe, Z. He, A. Kubica, and E. Tang, Layer codes as partially self-correcting quantum memories, *arXiv:2510.06659* (2025).
 - [18] T.-C. Lin, H.-P. Wang, and M.-H. Hsieh, Proposals for 3D self-correcting quantum memory, *arXiv:2411.03115* (2024).
 - [19] B. Roberts, J. M. Koh, Y. Tan, and N. Y. Yao, Cored product codes for quantum self-correction in three dimensions, *arXiv:2510.05479* (2025).
 - [20] R. Ma, V. Khemani, and S. Sang, Circuit-based characterization of finite-temperature quantum phases and self-correcting quantum memory, *arXiv:2509.15204* (2025).
 - [21] S. Chesi, B. Röthlisberger, and D. Loss, Self-correcting quantum memory in a thermal environment, *Phys. Rev. A* **82**, 022305 (2010).
 - [22] A. Al-Shimary, J. R. Wootton, and J. K. Pachos, Lifetime of topological quantum memories in thermal environment, *New J. Phys.* **15**, 025027 (2013).
 - [23] C. D. Freeman, C. M. Herdman, D. J. Gorman, and K. B. Whaley, Relaxation dynamics of the toric code in contact with a thermal reservoir: Finite-size scaling in a low-temperature regime, *Phys. Rev. B* **90**, 134302 (2014).

- [24] F. L. Pedrocchi, S. Chesi, and D. Loss, Quantum memory coupled to cavity modes, *Phys. Rev. B* **83**, 115415 (2011).
- [25] F. L. Pedrocchi, A. Hutter, J. R. Wootton, and D. Loss, Enhanced thermal stability of the toric code through coupling to a bosonic bath, *Phys. Rev. A* **88**, 062313 (2013).
- [26] A. Hutter and D. Loss, Breakdown of surface-code error correction due to coupling to a bosonic bath, *Phys. Rev. A* **89**, 042334 (2014).
- [27] S. Diehl, E. Rico, M. A. Baranov, and P. Zoller, Topology by dissipation in atomic quantum wires, *Nat. Phys.* **7**, 971 (2011).
- [28] C.-E. Bardyn and T. Karzig, Exponential lifetime improvement in topological quantum memories, *Phys. Rev. B* **94**, 094303 (2016).
- [29] C. D. Freeman, C. M. Herdman, and K. B. Whaley, Engineering autonomous error correction in stabilizer codes at finite temperature, *Phys. Rev. A* **96**, 012311 (2017).
- [30] S. Kumar and H. Weimer, Self-correction phase transition in the dissipative toric code, *arXiv:2602.19288* (2026).
- [31] G. D nnweber, G. Styliaris, and R. Trivedi, Quantum memory and autonomous computation in two dimensions, *arXiv:2601.20818* (2026).
- [32] Y.-J. Liu and S. Lieu, Dissipative phase transitions and passive error correction, *Phys. Rev. A* **109**, 022422 (2024).
- [33] G. S. Uhrig, Keeping a quantum bit alive by optimized π -pulse sequences, *Phys. Rev. Lett.* **98**, 100504 (2007).
- [34] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, Optimized dynamical decoupling in a model quantum memory, *Nature* **458**, 996 (2009).
- [35] H.-P. Breuer, E.-M. Laine, J. Piilo, and B. Vacchini, Colloquium: Non-markovian dynamics in open quantum systems, *Rev. Mod. Phys.* **88**, 021002 (2016).
- [36] A. Rivas, S. F. Huelga, and M. B. Plenio, Quantum non-markovianity: characterization, quantification and detection, *Rep. Prog. Phys.* **77**, 094001 (2014).
- [37] H.-P. Breuer, E.-M. Laine, and J. Piilo, Measure for the degree of non-markovian behavior of quantum processes in open systems, *Phys. Rev. Lett.* **103**, 210401 (2009).
- [38] A. Rivas, S. F. Huelga, and M. B. Plenio, Entanglement and non-markovianity of quantum evolutions, *Phys. Rev. Lett.* **105**, 050403 (2010).
- [39] D. Tamascelli, A. Smirne, S. F. Huelga, and M. B. Plenio, Nonperturbative treatment of non-markovian dynamics of open quantum systems, *Phys. Rev. Lett.* **120**, 030402 (2018).
- [40] M. Xu, J. T. Stockburger, and J. Ankerhold, Markovian embeddings of non-Markovian open system dynamics, *arXiv:2602.21430* (2026).
- [41] E. B. Davies, Markovian master equations, *Commun. Math. Phys.* **39**, 91 (1974).
- [42] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* **21**, 1087 (1953).
- [43] W. K. Hastings, Monte carlo sampling methods using markov chains and their applications, *Biometrika* **57**, 97 (1970).
- [44] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, A new algorithm for monte carlo simulation of ising spin systems, *J. Comput. Phys.* **17**, 10 (1975).
- [45] G. E. Uhlenbeck and L. S. Ornstein, On the theory of the brownian motion, *Phys. Rev.* **36**, 823 (1930).
- [46] T. Bergamaschi, R. Gheissari, and Y. Liu, Rapid mixing for gibbs states within a logical sector: a dynamical view of self-correcting quantum memories, *arXiv:2507.10976* (2025).
- [47] Y. Hong, J. Guo, and A. Lucas, Quantum memory at nonzero temperature in a thermodynamically trivial system, *arXiv:2403.10599* (2024).
- [48] N. Baspin, The free energy barrier: An Eyring–Polanyi bound for stabilizer Hamiltonians, with applications to quantum error correction, *arXiv:2509.17356* (2025).
- [49] E. Novais, E. R. Mucciolo, and H. U. Baranger, Resilient quantum computation in correlated environments: A quantum phase transition perspective, *Phys. Rev. Lett.* **98**, 040501 (2007).
- [50] J. Garcia Nila and T. A. Brun, Continuous quantum correction on Markovian and non-Markovian models, *arXiv:2505.18400* (2025).
- [51] R. Nandkishore and D. A. Huse, Many-body localization and thermalization in quantum statistical mechanics, *Annu. Rev. Condens. Matter Phys.* **6**, 15 (2015).
- [52] B. Bauer and C. Nayak, Area laws in a many-body localized state and its implications for topological order, *J. Stat. Mech.*, P09005 (2013).
- [53] G. Canossa, L. Pollet, M. A. Martin-Delgado, H. Song, and K. Liu, Error resilience of fracton codes and near saturation of code-capacity threshold in 3D, *arXiv:2512.22888* (2025).
- [54] C. M. Hopkin, V. V. Albert, and D. J. Williamson, Translation-invariant quantum low-density parity-check codes from compactified fracton models, *arXiv:2605.19298* (2026).
- [55] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2002).
- [56] D. R. Cox and D. Oakes, *Analysis of Survival Data* (Chapman and Hall, London, 1984).
- [57] E. T. Jaynes and F. W. Cummings, Comparison of quantum and semiclassical radiation theories with application to the beam maser, *Proc. IEEE* **51**, 89 (1963).
- [58] J. G. Kirkwood, Statistical mechanics of fluid mixtures, *J. Chem. Phys.* **3**, 300 (1935).
- [59] P. Talkner and P. Hanggi, Colloquium: Statistical mechanics and thermodynamics at strong coupling: Quantum and classical, *Rev. Mod. Phys.* **92**, 041002 (2020).
- [60] Particle Data Group, Review of particle physics, *Phys. Rev. D* **110**, 030001 (2024).
- [61] W. K. Grassmann, Transient solutions in markovian queueing systems, *Comput. Oper. Res.* **4**, 47 (1977).
- [62] D. T. Gillespie, Exact numerical simulation of the ornstein–uhlenbeck process and its integral, *Phys. Rev. E* **54**, 2084 (1996).