Self-correction in Wegner's three-dimensional Ising lattice gauge theory

David Poulin,[1,2,*] Roger G. Melko,[3,4,†] and Matthew B. Hastings[5,6,‡]
[1]*Département de Physique & Institut Quantique, Université de Sherbrooke, Québec, Canada J1K 2R1*
[2]*Canadian Institute for Advanced Research, Toronto, Ontario, Canada M5G 1Z8*
[3]*Department of Physics and Astronomy, University of Waterloo, Ontario, Canada N2L 3G1*
[4]*Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada N2L 2Y5*
[5]*Station Q, Microsoft Research, Santa Barbara, California 93106-6105, USA*
[6]*Quantum Architecture and Computation Group, Microsoft Research, Redmond, Washington 98052, USA*

(Received 17 December 2018; revised manuscript received 18 February 2019; published 8 March 2019)

Motivated by the growing interest in self-correcting quantum memories, we study the feasibility of self-correction in classical lattice systems composed of bounded degrees of freedom with local interactions. We argue that self-correction, including a requirement of stability against external perturbation, cannot be realized in system with broken global symmetries such as the two-dimensional Ising model, but that systems with local, i.e., gauge symmetries, have the required properties. Previous work identified a three-dimensional (3D) *quantum* system which realizes a self-correcting *classical* memory [E. Dennis *et al.*, J. Math. Phys. **43**, 4452 (2002)]. Here, we show that a purely classical 3D system, Wegner's Ising lattice gauge model [F. J. Wegner, J. Math. Phys. **12**, 2259 (1971)], can also realize this self-correction despite having an extensive ground-state degeneracy. We give a detailed numerical study to support the existence of a self-correcting phase in this system, even when the gauge symmetry is explicitly broken. More generally, our results obtained by studying the memory lifetime of the system are in quantitative agreement with the phase diagram obtained from conventional analysis of the system's specific heat [I. S. Tupitsyn *et al.*, Phys. Rev. B **82**, 085114 (2010)], except that self-correction extends beyond the topological phase, past the lower critical temperature.

DOI: 10.1103/PhysRevB.99.094103

I. INTRODUCTION AND MOTIVATION

A self-correcting memory is a passive physical device that stores information robustly at finite temperature despite fluctuations of its external parameters like magnetic field, pressure, etc. Such a system can be prepared in one of a finite number of initial states. The identity of this initial state can in principle be reconstructed with high probability up to some mixing time, after which all signatures of the initial configuration are lost. We say that a system is self-correcting if the mixing time grows with the system size. In particular, we are interested in self-correcting memories that arise from systems composed of localized, bounded degrees of freedom with short-range interactions.

It is well established that reliable computation can be realized from unreliable components using fault-tolerant techniques in both the classical [1–4] and quantum [5–11] settings. Moreover, fault-tolerant computation can be realized from local cellular automaton [12–14], which in physical terms corresponds to a lattice system with local interactions that are periodic in time. However, these processes dissipate heat. In contrast, a self-correcting memory is stabilized thermodynamically, i.e., through its interaction with a heat reservoir, and does not require external power.

To serve as a memory, a system must possess more than one metastable state, i.e., it must display phase coexistence. Information is stored in the system by preparing one of these states, and the probability of information corruption, i.e., the probability that the system spontaneously transitions from one phase to another, is exponentially suppressed with increasing system size.

Robust phase coexistence is however ruled out by the Gibbs phase rule [15]: for a system with N external parameters, the coexistence of P stable phases can only occur in a submanifold of codimension $P-1$. In other words, phase coexistence requires fine tuning of the system's parameters. As we will explain below, this rule can easily be understood in the Landau-Ginzburg paradigm of local order parameters and spontaneously broken symmetries, where it also rules out the existence of self-correcting memories. In this paper, we show how the Gibbs phase rule can be circumvented by turning to systems with nonlocal order parameters.

The motivations to study this problem stem from many sources. In the classical setting, the question of emerging global order from local noisy interactions has a long history. For instance, the phase rule forbids the existence of a phase transition in a one-dimensional, translationally invariant system composed of bounded, discrete degrees of freedom. However, Gács' cellular automaton [16,17] provides a "counterexample" to this rule, that is made possible using a time-dependent periodic Hamiltonian instead of a constant Hamiltonian [18]. Here, we present an alternative way of escaping the phase rule.

*david.poulin@usherbrooke.ca
†rgmelko@uwaterloo.ca
‡mahastin@microsoft.com

In the quantum setting, the prospect of quantum technologies has generated a growing interest in robust quantum memories. While the theory of fault-tolerant quantum computation is well developed, a self-correcting quantum memory [19] would enable passive, reliable quantum information storage, analogous to classical hard drives. The quantum model of [10,20] realizes a self-correcting quantum memory in four spatial dimensions and a classical memory in three spatial dimensions. It leaves open, however, the questions of a robust classical memory in a classical system, and of a robust quantum memory in lower spatial dimensions.

Lower-dimensional quantum model systems have been proposed [19,21–25] and disputed [26–28]. Thus, the existence of a self-correcting quantum memory in less than four spatial dimensions remains an open question to date. Given this status, it seams reasonable to step back and study classical self-correcting memories in a classical system before turning to the quantum setting. In particular, the study of self-correcting quantum memories in two spatial dimensions [19,21–23] appears premature given that our construction of a classical self-correcting memory is the only model we are aware of and requires three spatial dimensions.

In addition, the conditions for a quantum memory are more stringent than those for a classical memory. Suppose the energies of two memory states 0 and 1 differ by some unknown, possibly fluctuating, quantity Δ. If Δ is extensive, i.e., if there is a constant energy density difference between the two states, then thermal fluctuations will drive the memory to the state of lower energy. This thermal instability is a problem for both quantum and classical memories. However, a quantum memory must not only preserve the discrete memory states 0 and 1, it must also preserve coherent superpositions thereof $|\psi(0)\rangle = \alpha|0\rangle + \beta|1\rangle$. For any finite energy splitting Δ between the two states, Schrödinger's time evolution will introduce an unknown phase over time $|\psi(t)\rangle = \alpha|0\rangle + e^{-i\Delta t/\hbar}\beta|1\rangle$, effectively destroying the quantum superposition. Such a dephasing problem only affects quantum memories and does not require an extensive energy difference Δ, showing that the criteria to realize a self-correcting quantum memory are much more stringent than for a classical memory. In contrast, quantum mechanical models are richer than classical models in some fixed spatial dimension, so the question of classical self-correction in classical systems is a nontrivial intermediate case to study.

The rest of this paper is organized as follows. In the next section, we argue that systems with a broken local symmetry cannot realize a self-correcting memory. In Sec. III, we turn to systems with a gauge symmetry and explain how the previous argument breaks down. Section IV presents a detailed model and explains why it is stable against both external perturbations and thermal fluctuations. Numerical simulations of that model are presented in Sec. V where the self-correcting behavior is clearly observed. We conclude with general remarks and some open questions.

II. SYMMETRY-BROKEN PHASE

To build intuition, consider a two-dimensional ferromagnetic Ising model with Hamiltonian $H = -J\sum_{\langle i,j\rangle}\sigma_i\sigma_j$, where spins take values $\sigma_j \in \{-1, +1\}$ (see, e.g., [29]). It has two ground states, consisting of either all spins up or all spins down, and this degeneracy can be used to store one bit of information. At nonzero temperature, thermal fluctuations will create droplets of inverted spins, reducing the magnetization. However, the energy of an error droplet D is proportional to the length of its boundary $E = 2J|\partial D|$, so at temperature T, large droplets are suppressed by a Boltzmann factor e^{-E/k_BT} (we set $k_B = 1$). Below the critical Curie temperature T_C, error droplets are typically too sparse to percolate, so the sign of the magnetization retains information about the stored bit of information. Statistical fluctuations could lead to spontaneous percolation of the droplets, but such fluctuations are exponentially suppressed with the system size, meaning that the memory lifetime grows exponentially with the system size.

This stability is attributable to a symmetry: flipping all the spins leaves the Hamiltonian invariant. Subcritical temperature thermal states spontaneously break this symmetry, either choosing a positive or negative magnetization sector, thus enabling the system to serve as a memory. This requires fine tuning, however: generic perturbations to the Hamiltonian will break this symmetry and favor one sector over the other. For instance, a magnetic field will add a term $B\sum_i \sigma_i$ to the Hamiltonian. The energy of a spin-down droplet D in a spin-up background is modified to $E = -2B|D| + 2J|\partial D|$. Because the area $|D|$ of a droplet grows faster than its boundary $|\partial D|$, no matter how small the magnetic field B is, the Boltzmann factor will favor the formation of large droplets, and hence a unique stable sector. We see that phase coexistence is restricted to a codimension 1 manifold of the (B, T) phase diagram in accordance to Gibbs' phase rule.

Because of the required fine tuning, the Ising model is not a robust phase of matter, it is merely a symmetry-protected phase of matter. A self-correcting memory is a symmetry-broken finite-temperature phase, robust to generic physical perturbations that are not constrained by any symmetries. This is by definition impossible to realize in symmetry-broken phases with a *local* order parameter. Suppose indeed that the Hamiltonian H has a spontaneously broken symmetry with associated order parameter M, i.e., M takes distinct expectation values in the different thermal sectors. Adding a perturbation μM to the Hamiltonian will favor one sector over the others, so the system will thermalize to this unique sector irrespective of initial conditions. Moreover, because M is a local order parameter, the term μM consists of physically realistic local interactions. In the Ising model, for instance, M would be magnetization and the perturbation would correspond to an external magnetic field.

III. NONLOCAL ORDER

We have argued above that in the presence of symmetry-breaking perturbations, thermal stability is incompatible with the existence of a local order parameter. Instead, we turn to systems with nonlocal order parameters. In the quantum setting, it is possible for a system to possess two distinct

ground states $|\psi_0\rangle$ and $|\psi_1\rangle$ that are locally indistinguishable. The two states,

$$|\psi_0^{AB}\rangle = \frac{1}{\sqrt{2}}(|\uparrow^A\uparrow^B\rangle + |\downarrow^A\downarrow^B\rangle), \tag{1}$$

$$|\psi_1^{AB}\rangle = \frac{1}{\sqrt{2}}(|\uparrow^A\uparrow^B\rangle - |\downarrow^A\downarrow^B\rangle), \tag{2}$$

illustrate this idea. In both, the single-spin state obtained from a partial trace is maximally mixed

$$\rho_{0/1}^A = \rho_{0/1}^B = \tfrac{1}{2}(|\uparrow\rangle\langle\uparrow| + |\downarrow\rangle\langle\downarrow|), \tag{3}$$

yet the two states are orthogonal to each other. Thus, any single-spin measurement cannot distinguish between the two states. Topologically ordered systems extend this idea to local Hamiltonians whose degenerate ground states are indistinguishable from each other given any observable acting on a homologically trivial region of the system manifold space. Thus, these systems naturally present a stability to local perturbations [30–32].

In the classical setting, two distinct configurations of a system composed of local degrees of freedom must unavoidably differ locally, so are subject to an energy splitting by a local field. To obtain a self-correcting memory, we therefore cannot encode information states in distinct system configurations. Instead, we choose a one-to-many encoding where each information state is encoded in an ensemble of classical states, e.g., 0 is encoded in any configuration from the ensemble $\Omega_0 = \{\sigma_0^1, \sigma_0^2, \ldots\}$ and 1 is encoded in any configuration from the ensemble $\Omega_1 = \{\sigma_1^1, \sigma_1^2, \ldots\}$. While any of these individual states can be locally distinguished, the distinct ensembles can be chosen to be statistically locally indistinguishable.

A concrete way of realizing this encoding uses a Hamiltonian with a gauge symmetry. The spectrum of such a model is exponentially degenerate because configurations related by a gauge transformation have the same energy. The information ensembles Ω_j will consist of sets of states that are related by a gauge transformation. The fact that these ensembles cannot be distinguished locally is then a corollary of Elitzur's theorem [33,34] which states that a gauge symmetry cannot be spontaneously broken. Moreover, the nonlocal order parameter characterizing the gauge model remains well defined in the absence of a gauge symmetry, so it can continue to serve as the information readout in the model even when a symmetry-breaking field is added. The next section illustrates these ideas with a concrete model, Wegner's three-dimensional (3D) Ising lattice gauge theory [35].

IV. MODEL

The model we consider is a cubic lattice of size $L \times L \times L$, with Ising spins $\sigma_i \in \{-1, +1\}$ residing on each *edge* (see Fig. 1 for an illustration). The Hamiltonian is the sum over plaquettes

$$H_0 = -J \sum_P A_P, \tag{4}$$

where $J > 0$ and for plaquette P with boundary edges $\partial P = (i, j, k, l)$, we define $A_P = \sigma_i\sigma_j\sigma_k\sigma_l$. For simplicity, we assume the lattice has periodic boundary conditions. Beyond its use in lattice gauge theory [34], this model is familiar in quantum error correction as it consists of the"classical part" of the three-dimensional toric code [10,20]. That quantum model is known to be a self-correcting classical memory; the question we ask here is whether it retains that property without the kinetic quantum terms. Note that the spectrum of the model changes drastically with the presence of the quantum kinetic term. In particular, the quantum model has some constant ground-state degeneracy while the degeneracy of the classical model is extensive. Thus, there is *a priori* no reason to assume any relation between the thermal properties of the two models, particularly in the presence of an external field as we explain below.

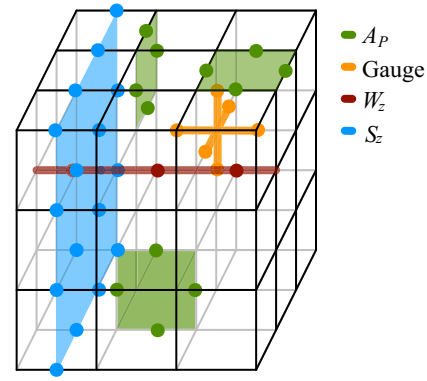

FIG. 1. The model consists of Ising spins σ_i placed on each *edge* of a regular square lattice with periodic boundary conditions. Green plaquettes represent the coupling terms of the Hamiltonian. Orange vertices represent gauge symmetries. The red string is a gauge-invariant Wilson loop W_z. Flipping all spins along the blue plane S_z perpendicular to the z axis inverts the value of W_z.

This Hamiltonian has a gauge symmetry generated by flipping all six spins adjacent to a vertex. To see that this is a symmetry, note that a plaquette and a vertex share an even number η (either 0 or 2) of edges, so the gauge transformation will change a plaquette coupling A_P by a multiplicative factor $(-1)^\eta = 1$. In other words, this follows from the fact that the boundary of a plaquette has itself no boundary, $\partial\partial P = 0$.

A. Observables

A gauge-invariant observable is obtained by considering the product of the spins along any closed loop ℓ, generating a so-called Wilson loop $W_\ell = \prod_{i\in\ell}\sigma_i$. The gauge invariance follows from $\partial\partial\ell = 0$. For loops ℓ that are the boundary of a surface S, i.e., $\ell = \partial S$, the corresponding Wilson loops are obtained by taking the product of the enclosed plaquette terms $W_{\partial S} = \prod_{P\in S} A_P$, so they are coupled to the energy density in S. High-temperature expansions can be used to show [34,35] that $\langle W_\ell\rangle \sim e^{-\alpha|S|}$ where α is some function of βJ and S is the minimal surface area enclosed by the loop, i.e., for which $\partial S = \ell$. At low temperature, however, any single spin flip along ℓ will invert the sign of W_ℓ, so the expectation of a Wilson loop vanishes exponentially with its length $\langle W_\ell\rangle \sim e^{-\gamma|\ell|}$ where to leading order $e^{-\gamma}$ is the probability of an unflipped spin, so $e^{-\gamma} \approx (1 - e^{-4J/T})$. This discrepancy between high- and low-temperature calculations suggests the existence of a finite-temperature phase transition. Indeed, this

model is dual to a 3D Ising model [34,35] whose critical temperature is well established numerically $T_c \approx 1.314J$ [36].

A system embedded on a topologically nontrivial manifold, such as a cube with periodic boundary conditions or a punctured cube, admits homologically nontrivial loops $\ell \neq \partial S$ to which the above high-temperature expansion does not apply. These Wilson loops are gauge-invariant observables that are decoupled from the energy. One such example, illustrated on Fig. 1, is the loop which winds around the three-dimensional torus in the z direction W_z.

To see that W_z is indeed decoupled from the energy, consider the dual plane S_z which consists of all the edges in the z direction with a fixed z coordinate (see Fig. 1). Inverting all the spins in S_z leaves all plaquette terms A_P invariant because S_z intersects A_P on an even number of edges. On the other hand, W_z and S_z intersect on an odd number edges, so flipping all spins in S_z inverts the value of W_z. We conclude that, given any spin configuration σ, there exists another configuration σ' obtained by flipping all the spins in S_z such that $H_0(\sigma) = H_0(\sigma')$ and $W_z(\sigma) = -W_z(\sigma')$. We encode one bit of information using the two ensembles of spin configurations Ω_+ and Ω_- consisting of the ground states of H_0 with $W_z = +1$ and $W_z = -1$, respectively.

B. Stability against perturbation

To understand the robustness of this phase, let us repeat the above argument in the presence of a generic perturbation $V = \sum_i v_i$ to the Hamiltonian $H = H_0 + V$, where each term v_i is a bounded function $|v_j| \leqslant K$ of all the spins within a constant radius r away from spin i. Given a spin configuration σ of energy $H(\sigma)$, consider the configuration σ' obtained by flipping the spins in S_z. It is easy to see that the two configurations have a vanishing energy density difference as the volume $\mathcal{V} = L^3$ grows:

$$\frac{|H(\sigma) - H(\sigma')|}{\mathcal{V}} = \frac{|V(\sigma) - V(\sigma')|}{\mathcal{V}} \leqslant \frac{4Kr}{\mathcal{V}^{1/3}}. \tag{5}$$

In the thermodynamic limit, thermal fluctuations will thus not discriminate between the two sectors, so there is phase coexistence. More strongly, the local density matrices of the two different ensembles agree up to exponentially small error on any finite-size patch, so that the energy difference is exponentially small in L.

C. Thermal stability

Thermal stability is more subtle. Suppose we prepare a ground state σ of H_0 with a fixed value of $W_z = w$ and let the system thermalize. In the ground state, we have $A_P = 1$ for all P, but in thermal equilibrium some of the A_P will take value -1. Excited states are obtained from the ground state by flipping spins contained in dual-membranes: two-dimensional submanifolds akin of the region S_z but which do not span an entire plane. Just like in the two-dimensional Ising model, the energy of such an error membrane M grows proportionally to its boundary $E = 2J|\partial M|$, so below a critical temperature the membranes are confined. Despite this confinement, however, at any nonzero temperature, we expect a constant density of membranes. Since the sign of W_z is inverted every time it intersects an error membrane, W_z averages to 0 in the thermodynamic limit as $1 - e^{-\gamma L_z}$ where $e^{-\gamma}$ is roughly the density of error membranes.

Despite the vanishing of W_z, the system retains some information about its initial configuration. After letting the system interact with a heat bath for some time, suppose we were to cool it down. Error membranes would slowly shrink and disappear, returning the system to a spin configuration with the same value $W_z = w$ as initially, unless the thermalizing and cooling processes have generated membranes that fused into a large membrane that spans a homologically nontrivial plane. Because error membranes are confined, the probability of such an event vanishes in the thermodynamic limit.

While cooling the system back to $T = 0$ is not physically possible, it shows that, in principle, the information is still present in the system. In fact, the cooling phase does not need to be implemented to read out the information. It is used here only to illustrate that the information is still encoded in the system, although in a hidden form. There exists an alternative "algorithmic" procedure for retrieving the information. It consists of reading out the thermal spin configuration σ entirely, and recording the value of each individual spin. From this knowledge, the value of each A_P can be computed. The P with $A_P = -1$ indicate the boundary ℓ of error membranes. Then, an algorithmic procedure called *decoding* can be used to determine what are the smallest membranes consistent with these boundaries: $\arg\min_M\{|M| : \partial M = \ell\}$. This may be computationally expensive but can certainly be done in principle. Inverting the spins contained in these minimal membranes should return the system to its initial value of W_z unless some error membranes have grown and percolated to reach a macroscopic size. Thus, the self-correcting phase truly corresponds to a membrane confinement phase. Thus, we expect the self-correcting phase diagram to coincide with the one obtained for confinement, which is dual to the 3D Ising model with $h = 0$.

While measuring all the spins is impractical for a real material, the protocol presented above could be of practical use in a metamaterial, where spins are engineered systems such as solid-state devices. For real materials, even if the procedure is impractical, it does illustrate that the information has not been irreversibly erased from the system.

D. General stability

We consider thermal stability in the presence of a gauge symmetry-breaking field. Suppose a magnetic field $V = -B\sum_j \sigma_j$ is added, energetically favoring $\sigma_j = +1$. The unique ground state σ^0 of $H = H_0 + V$ is the all spin-up configuration, and has $W_z = +1$. The minimal energy configuration σ^1 with $W_z = -1$ is obtained from σ^0 by flipping all spins in S_z. Error membranes *restricted to the plane* S_z have the same energetic cost as those in the two-dimensional Ising model $E = -2B|M| + 2J|\partial M|$. At first glance, the same argument invoked for the Ising model suggests that large error membranes will proliferate in this plane and result in a configuration in the $W_z = +1$ sector, corrupting the memory.

However, at nonzero temperature, entropy will make the error membranes fluctuate in and out of the S_z plane. Outside the plane, the magnetic field contributes positively to the energy. There is a critical temperature where the entropy gained from

fluctuating the error membrane in and out of S_z compensates for the energy gained from restricting the error membrane to S_z. Thus, the system does not magnetize even in the presence of an external field, which is again just a restatement of Elitzur's theorem. The system with an external field is self-dual and has been studied numerically in [37]. We expect the self-correcting upper-temperature phase transition to coincide with the topological phase reported in this reference. What is less clear is whether a lower critical temperature is required to prevent large error droplets from proliferating in the S_z plane. We further investigate this question in Sec. V C.

V. NUMERICAL RESULTS

A. Setting

We verified the above claims numerically. The physical process we are simulating is the following. The system is prepared at some low temperature T_0 with all spins up, except in a single S_z plane where all spins are down, so initially $W_z = -1$. With a positive magnetic field B, this initial configuration is not a ground state because the magnetic field favors up spins. The temperature is then slowly ramped up from T_0 to some holding temperature T_{hold}. The system sits at this holding temperature for some finite amount of time t, after which the temperature is ramped back down to T_0. At that time, we measure the Wilson loop W_z to check if it retained information about its initial value. Note that we do not need to run separate experiments for the two initial values of $W_z = \pm 1$ since the external field favors $W_z = +1$, i.e., the noise affecting the encoded information is asymmetric. The challenging case is when the system carries the information $W_z = -1$, and so we focus on that case.

We use standard Monte Carlo simulations with Metropolis-Hastings rule. A single Monte Carlo update implements both single-spin flips and cluster updates corresponding to the gauge symmetry. The conclusions we reach are independent of the specific choice of updates. Each update consists of N_s attempted single-spin flips, where N_s is the number of spins, and N_v gauge cluster updates, where N_v is the number of lattice vertices. This way, the number of Monte Carlo updates corresponds to a physical measure of time, independent of the lattice size. We fix $T_0 = 0.2J$. Ramping the temperature is performed with steps of $\Delta T = 0.05J$. We use 5000 Monte Carlo updates at each temperature. The system is held at T_{hold} for 10 000 Monte Carlo sweeps, before cooling back down to T_0. Results below are averaged over 1000 such temperature sweeps.

B. Self-correction

Simulation results are shown in Fig. 2 where we report the final expectation value of the Wilson loop as a function of the holding temperature for various system sizes and external fields. We emphasize that the main plot is not the expectation values of the Wilson loop as a function of temperature during a ramping process. Each data point is measured at T_0, but the measurement is preceded by a temperature ramp-up to some holding temperature followed by a ramp-down. The data are then plotted as a function of the holding temperature. An example of a single temperature sweep is illustrated in the inset for $B = 0$. The inset shows strongly polarized values of $\langle W_z \rangle$, in apparent conflict with our claim that should vanish as $\langle W_z \rangle \approx e^{-\gamma L}$ at any finite temperature. However, this is due to finite-size effects since $\gamma \approx -\log(1 - e^{-4J/T}) \approx 0.047$ is quite small.

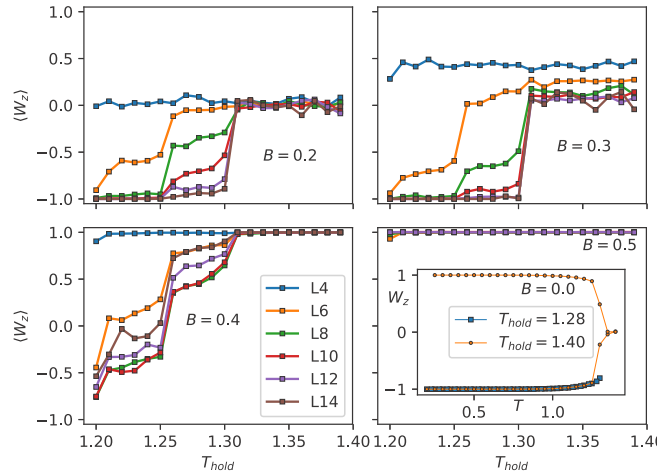

FIG. 2. Expectation value of the Wilson loop $\langle W_z \rangle$ as a function of the holding temperature T_{hold} for various system sizes and external magnetic fields B. For sufficiently weak external fields and sufficiently low holding temperatures, the system returns to its initial value $W_z = -1$ with a probability that increases with system size L, suggesting the existence of a self-correcting phase in the (T_{hold}, B) diagram. At $B = 0.4J$, we observe that the system of size $L = 12$ has a higher error probability than $L = 10$ for all holding temperatures, suggesting that this B exceeds the critical field value. The inset illustrates one typical temperature sweep for two different holding temperatures for a system of size $L = 12$. Each data point in the main plots correspond to an average of 1000 such temperature sweeps.

We say that the system with parameters (T_{hold}, B, J) is in the self-correcting phase when the probability that the system returns to a different value of W_z decreases exponentially with the system size L. We can verify this claim by fitting $\langle W_z \rangle$ as a function of the system size in the subcritical region, which is shown in Fig. 3.

Based on these results, we predict the existence of an upper critical temperature $T_c \approx 1.31J$ and an upper critical field $0.3J \lesssim B_c \lesssim 0.4J$ such that the system is self-correcting

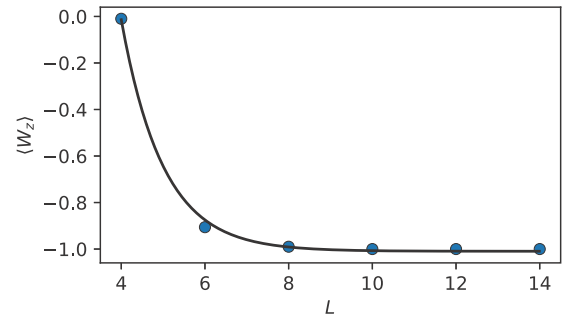

FIG. 3. Final Wilson loop expectation as a function of the system size for subcritical holding temperature $T_{\text{hold}} = 1.2J < T_c$ and external field $B = 0.2J < B_c$. The result agrees with a fit $\langle W_z \rangle = a \exp(-bL) - 1$, shown in the solid line.

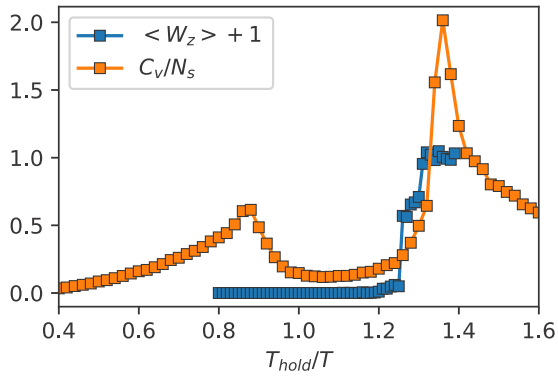

FIG. 4. A comparison of the final Wilson loop expectation value for $B = 0.2J$ and $L = 8$ (compare to Fig. 2) and the specific heat per spin, calculated as a typical thermodynamic estimator, for the same parameter values.

for $T_{\text{hold}} < T_c$ and $B < B_c$. The critical temperature is largely insensitive to the magnetic field until a paramagnetic phase is obtained. These observations are consistent with the well-established critical temperature $T_c \approx 1.314J$ at $B = 0$ [36] and more generally with the findings of [37] which shows an upper critical field at $B_c \approx 0.225T$ which is $0.29J$ at $T = T_c$, in agreement with our observation.

C. Lower critical temperature?

The 3D toric code is known to be a self-correcting quantum memory [10,20]. Its Hamiltonian $H_Q = H + K$ is the sum of Wegner' classical gauge model, Eq. (4), and a quantum kinetic term K that commutes with H. Because of this commutation, the partition function factors $Z_\beta(H_Q) = Z_\beta(H) \times Z_\beta(K)$ so, in this case, the relations between the thermal properties of the quantum and classical model are expected. But, in the presence of a magnetic field, this factorization breaks and the thermal stability of the classical model is not a direct consequence of the stability of the quantum model.

In particular, the quantum model is protected from an external magnetic field by the kinetic term, which gives it a spectral gap [30–32]. The classical model is only protected by entropy, an effect that we may term "topological order by disorder." For this reason, and as discussed in Sec. IV D, we expect stability only at nonzero temperature. If we consider, for instance, the case of $B = 0.2J$, the phase diagram proposed in [37] predicts a lower critical temperature of $T_\ell \approx 0.89J$. However, we do not observe this lower critical temperature when simulating the system's ability to self-correct. Figure 4 shows the expectation value of the Wilson loop as a function of the holding temperatures and the specific heat as a function of temperature. While the phase transition is clearly visible on the specific head and coincides with the predictions of [37], there is no sign of a transition in the self-correction data.

This discrepancy is not necessarily a problem. The phase diagram represents the equilibrium properties of a system, while self-correction is truly about the equilibration process itself. It is possible that the equilibration process remains exponentially slow below the lower critical temperature. It is also possibly that the character of the equilibration process changes at this transition, perhaps transitioning from an exponential lifetime to a polynomial lifetime, which is difficult to distinguish in the system sizes we have simulated. In any case, it would be premature based on these simulations to conclude that the system remains self-correcting all the way to $T = 0$, and this question deserves further scrutiny.

VI. CONCLUSION

We have argued that a self-correcting memory can be realized in a lattice system with bounded degrees of freedom and local interactions if the Hamiltonian is endowed with a gauge symmetry and is embedded on a topologically nontrivial manifold. The gauge symmetry implies that the Hamiltonian spectrum is exponentially degenerate. The nontrivial topology gives rise to additional degeneracies that can be probed by homologically nontrivial Wilson loops. Information can be reliably stored in this Wilson loop and Elitzur's theorem tells us that the stored information will not respond to an external field.

We have tested these predictions numerically on Wegner's 3D Ising lattice gauge theory with an external magnetic field. Our results are consistent with the existence of a phase in the temperature/field diagram where the information is exponentially well retained as a function of the system size, in agreement with the model's previously established phase diagram.

The argument for the stability of the model to external fields relies on entropy, so is only expected to hold for a nonzero temperature. A lower critical temperature has been observed in previous conventional characterizations of the model, but is not seen in our numerical study. This leaves open the question of whether temperature is necessary to stabilize the encoded information.

ACKNOWLEDGMENTS

We thank C. Bennett for stimulating discussions. This work is financially supported by NSERC, the Perimeter Institute for Theoretical Physics, and the Shared Hierarchical Academic Research Computing Network (SHARCNET). Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation.

[1] J. von Neumann, *Automata Studies* (Princeton University Press, Princeton, NJ, 1956), pp. 329–378.

[2] R. L. Dobrushin and S. I. Ortyukov, Problems Inf. Trans. **13**, 201 (1977).

[3] N. Pippenger, in *Proceedings of the 26th IEEE FOCS Symposium, Portland, USA, 1985* (IEEE, Piscataway, NJ, 1985), p. 30.

[4] D. Spielman, in *Proceedings of the 37th IEEE FOCS Symposium, Burlington, VT, USA, 1996* (IEEE, Piscataway, NJ, 1996), p. 154.

[5] P. W. Shor, in *Proceedings of the 37th IEEE FOCS Symposium, Burlington, VT, USA, 1996* (IEEE, Los Alamitos, California, 1996), pp. 56–65.

[6] A. M. Steane, Phys. Rev. A **54**, 4741 (1996).
[7] D. Aharonov and M. Ben-Or, in *Proceedings of the 29th Annual ACM Symposium on Theory of Computing, El Paso, Texas, USA, 1997* (ACM, New York, 1997); longer version arXiv:quant-ph/9906129, arXiv:quant-ph/9611025.
[8] J. Preskill, Proc. R. Soc. London, Ser. A **454**, 385 (1998).
[9] E. Knill, R. Laflamme, and W. H. Zurek, Science **279**, 342 (1998).
[10] A. Y. Kitaev, Ann. Phys. (NY) **303**, 2 (2003).
[11] E. Knill, Nature (London) **434**, 39 (2005).
[12] A. Toom, in *Multicomponent Systems*, edited by R. L. Dobrushin, Advanced in Probability, Vol. 6 (Dekker, New York, 1980), p. 549.
[13] E. R. Berlekamp, J. H. Conway, and R. K. Guy, *Winning Ways for your Mathematicl Phays* (Academic, New York, 1982).
[14] P. Gács, J. Comput. Syst. Sci. **36**, 125 (1988).
[15] J. W. Gibbs, *The Scientific Papers* (Dover, New York, 1961).
[16] P. Gács, J. Comput. Syst. Sci. **32**, 15 (1986).
[17] P. Gács, J. Stat. Phys. **103**, 45 (2001).
[18] C. H. Bennett and G. Grinstein, Phys. Rev. Lett. **55**, 657 (1985).
[19] D. Bacon, Phys. Rev. A **73**, 012340 (2006).
[20] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, J. Math. Phys. **43**, 4452 (2002).
[21] A. Hamma, C. Castelnovo, and Chammon, Phys. Rev. B **79**, 245122 (2009).
[22] A. Hutter, J. R. Wootton, B. Röthlisberger, and D. Loss, Phys. Rev. A **86**, 052340 (2012).
[23] F. Pedrocchi, A. Hutter, J. Wootton, and D. Loss, Phys. Rev. A **90**, 012321 (2014).
[24] K. P. Michnicki, Phys. Rev. Lett. **113**, 130501 (2014).
[25] C. G. Brell, New J. Phys. **18**, 013050 (2016).
[26] J. Haah and J. Preskill, Phys. Rev. A **86**, 032308 (2012).
[27] O. Landon-Cardinal and D. Poulin, Phys. Rev. Lett. **110**, 090502 (2013).
[28] O. Landon-Cardinal, B. Yoshida, D. Poulin, and J. Preskill, Phys. Rev. A **91**, 032303 (2015).
[29] M. Plischke and B. Bergersen, *Equilibrium Statistical Physics* (World Scientific, Singapore, 1994)
[30] S. Bravyi, M. B. Hastings, and S. Michalakis, J. Math. Phys. **51**, 093512 (2010).
[31] I. Klich, Ann. Phys. (NY) **325**, 2120 (2010).
[32] S. Bravyi and M. Hastings, Commun. Math. Phys. **307**, 609 (2011).
[33] S. Elitzur, Phys. Rev. D **12**, 3978 (1975).
[34] J. B. Kogut, Rev. Mod. Phys. **51**, 659 (1979).
[35] F. Wegner, J. Math. Phys. **12**, 2259 (1971).
[36] R. Gupta and P. Tamayo, US-Japan Bilateral Seminars, Maui, 1996 (unpublished).
[37] I. S. Tupitsyn, A. Kitaev, N. V. Prokof'ev, and P. C. E. Stamp, Phys. Rev. B **82**, 085114 (2010).