Self-Correcting Quantum Memories Beyond the Percolation Threshold

Matthew B. Hastings,[1,2] Grant H. Watson,[3] and Roger G. Melko[3,4]
[1]Microsoft Research, Station Q, CNSI Building, University of California, Santa Barbara, California 93106, USA
[2]Quantum Architectures and Computation Group, Microsoft Research, Redmond, Washington 98052, USA
[3]Department of Physics and Astronomy, University of Waterloo, Ontario N2L 3G1, Canada
[4]Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada
(Received 12 September 2013; published 20 February 2014)

We analyze several high dimensional generalizations of the toric code at a nonzero temperature. We find that in a large enough dimension, there can be a distinct separation between the critical temperature T_c, given by thermodynamic singularities, and the percolation temperature T_p, given by the percolation of defects. We argue that the regime $T_p < T < T_c$ is a range of temperatures where a self-correcting quantum memory can operate despite having percolating defects. We present analytic arguments and numerical evidence in support of this scenario, including a mean-field treatment and Monte Carlo simulations. Near T_c, simulations observe a large hysteretic behavior, which may have applications by allowing the self-correcting phase to survive in a "superheated" regime.

DOI: 10.1103/PhysRevLett.112.070501 PACS numbers: 03.67.Lx, 05.30.Pr, 64.60.ah

The classical Ising model is the prototypical example of a "self-correcting" memory. In two or more dimensions, in the low temperature ferromagnetic phase, the system can store a single classical bit for an exponentially long time in the sign of the global magnetization. In contrast, many proposals to protect quantum information against noise require active error correction by an external classical control [1]. Other proposals like low-dimensional topological quantum memories must avoid thermally excitating anyons [2], requiring a temperature that tends to zero as the inverse logarithm of the system size.

Remarkably, a topological memory in four dimensions (4D) [3] self-corrects up to a fixed nonzero temperature T_c, describing a transition between a high-temperature disordered phase, and a low-temperature phase capable of topologically encoding the quantum state. There have been extensive searches for lower-dimensional self-correcting memories [4], including the cubic code [5], though it still does not have a lifetime diverging arbitrarily largely with system size at $T > 0$ [6]. The question of true self-correction in $D < 4$ is still an outstanding open problem. However, it has recently been proposed that higher-dimensional toric "surface" codes could be artificially constructed [7], e.g., by building long-range connections between local components in a superconducting quantum circuit, motivating us to analyze in more detail the properties of higher dimensional models.

Conventional approaches use the percolation of thermally activated defects to estimate the T_c where the self-correcting phase is lost [8]. Surprisingly, using a combination of numerical simulation and gauge-invariant mean-field theory, we find that the phase transition temperature T_c of some higher dimensional toric codes is *not* the same as the temperature T_p for percolating defects, and that error correction is possible even when defects percolate. We show analytically that the ratio T_c/T_p can diverge as D becomes large, allowing error correction over an unexpectedly large range of temperatures. Further, Monte Carlo simulations find a large region of hysteresis, where the low-T phase is metastable well above T_c.

(p, q) Toric codes.—We begin by briefly describing the formalism for generalizations of the toric code to a hypercubic lattice in D dimensions. We refer to the different codes as (p, q) codes, with $p + q$ equal to the spatial dimension D. In this notation, the original toric code is a (1,1) code while the 4D self-correcting code is a (2,2) code. Following the usage in topology, we refer to the vertices of the lattices as 0-cells, the edges of the lattice as 1-cells, the plaquettes as 2-cells, and so on, up to D-cells. We use N_k to denote the number of k-cells. Note that $N_k = N_{D-k}$; this equality is related to a duality between (p, q) codes and (q, p) codes. In a (p, q) code, there is one spin 1/2 (or qubit) on each p-cell. The Hamiltonian is $H = J_A H_A + J_B H_B$, with J_A, J_B being positive scalars and

$$H_A = -\sum_{c_{p+1}} \prod_{i \in c_{p+1}} S_i^z, \qquad H_B = -\sum_{c_{p-1}} \prod_{i \ni c_{p-1}} S_i^x. \quad (1)$$

In H_A, the sum is over $p+1$-cells, denoted c_{p+1}, and the product is over p-cells, labeled by i, with $i \in c_{p+1}$ meaning that i is attached to c_{p+1}. The notation in H_B is similar, with the sum being over $p-1$-cells, c_{p-1}.

The partition function $Z(\beta) = \mathrm{tr}[\exp(-\beta H)]$ decomposes exactly as

$$Z(\beta) = Z_A(\beta) Z_B(\beta) 2^{-N_p}, \quad (2)$$

where $Z_A = \mathrm{tr}[\exp(-\beta H_A)]$ and $Z_B = \mathrm{tr}[\exp(-\beta H_B)]$. Hence, we can compute $Z(\beta)$ by calculating Z_A and Z_B separately, e.g., using classical Monte Carlo calculations.

In this Letter, we focus on $(1, D-1)$ codes to simplify the numerics. In fact, these codes are *not* self-correcting quantum memories—they can self-correct against spin flip errors but not against dephasing errors (S^z errors). A $(2, D-2)$ code can self-correct against both types of errors for $D \geq 4$ [9], and we expect it to have similar large D behavior. The self-dual code is the $(D/2, D/2)$ code—it may have very different large D behavior which we will address in future. In this Letter, we only simulate H_A, where the spins are on 1-cells (edges) and the interactions are on 2-cells (plaquettes); i.e., we study self-correction against spin flip errors only.

The (0, D) code is the Ising model, with no ability to correct against dephasing errors, being only useful as a classical memory. Here, the phenomenon that the percolation and phase transition temperatures are distinct is well established. It has been proven that $T_p \sim D/\log(D)$ for large D [10] while $T_c \sim D$ in agreement with mean-field theory (MFT) [11], and so for sufficiently large D, $T_p < T_c$. Further, numerical simulations [12] show that this occurs already in 3D. This difference between T_p and T_c means that the Peierls argument [13] cannot correctly predict T_c for the Ising model.

Defects, percolation temperature, and relation to error correction.—In the Ising model, the low temperature expansion sums over domain walls, called Peierls contours, between up and down spins, with a weight dependent on the area of the contour. It is easier to understand these contours in a *dual* version of the Ising model, a (D, 0) code, where the spins are on the D-cells, and the interactions are on the $D-1$-cells between a pair of D-cells. A spin configuration assigns $+1$ or -1 to each D-cell and the set of interactions which are unsatisfied is the boundary of this configuration. Since the set of unsatisfied interactions is a boundary, and the boundary of a boundary vanishes, the set of unsatisfied interactions indeed gives closed surfaces. For these surfaces made of $D-1$-cells, we regard two $D-1$-cells as being neighbors if they both attach to the same $D-2$-cell. For the (0, D) code, the defects are 1-cells instead, and two 1-cells are neighbors if they both attach to the same 2-cell.

For a $(1, D-1)$ code, the defect surfaces are sets of 2-cells, with two 2-cells being neighbors if they both attach to the same 3-cell (see Fig. 1). In a dual picture, defect surfaces are now closed $D-2$-dimensional surfaces.

Reference [10] upper bounds T_p by choosing a subset of Peierls contours which can be counted more easily. These contours are obtained by constructing a sequence of k spins starting from a given spin (say, at the origin of an infinite hypercubic lattice), flipping that spin, and then flipping each next spin in turn by shifting by distance 1 in any of the positive coordinate directions. For the given starting spin, there are D^{k-1} such sequences of spins giving entropy $\sim k\log(D)$, while the area of the contour is proportional to Dk. So for $\beta \lesssim \log(D)/D$, the energetic cost does not suppress the appearance of these chains. This bound generalizes straightforwardly to our problem, giving the same scaling $T_p \lesssim D/\log(D)$.

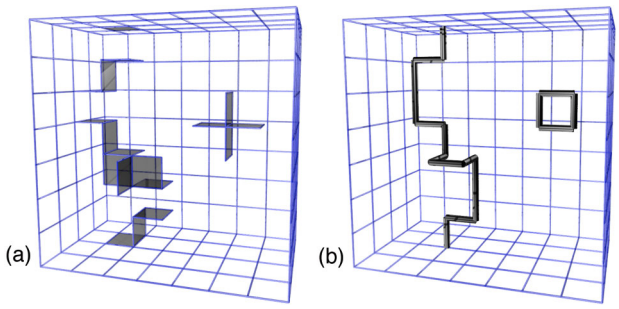

FIG. 1 (color online). Defect clusters in the 3D code. (a) A set of percolating (left) and nonpercolating (right) clusters of defect plaquettes in H_A, defined as being 2-cells that have an odd number of Ising variables S_i^z. (b) The defect surfaces in the dual picture.

We show below that for the $(1, D-1)$ code for large D, $T_c \sim D$, so $T_c > T_p$ for sufficiently large D. Thus, just as the Peierls argument cannot correctly predict T_c in the Ising model, arguments based on percolating defects cannot correctly predict T_c of certain high dimensional toric codes. However, we claim that T_c, rather than T_p, determines the upper temperature at which the code is a self-correcting memory. To act as a self-correcting memory, we need to define a recovery procedure. Then, we are interested in the procedure of encoding information into the memory at $T = 0$, heating the memory to some $T < T_c$ at which it stays for some time, and finally trying to recover the information (see Fig. 2). A recovery procedure involves measuring the set of defect plaquettes, called the "syndrome", and then applying spin flips to correct the defects before reading the encoded information. Since the spin flips arising from noise and from the recovery map one ground state to another, they are a 1 cocycle. If the 1 cocycle is topologically trivial, then the information can be recovered. For $T < T_c$, the dynamics are not critical and the system relaxes quickly; thus, in the thermodynamic limit, we expect that the cocycle will be topologically trivial. Numerical experiments below confirm that there indeed is a long lifetime in the low temperature phase; in these experiments, the system is heated from $T = 0$ to $T > 0$, then cooled back to $T = 0$ and finally the information is read. The cooling to $T = 0$ *defines* a recovery procedure by a randomized algorithm because, after returning to $T = 0$, we have returned to a state with no defects. Importantly, the Monte Carlo cooling back to $T = 0$ can be implemented only by measuring the syndrome to determine the spin flip probabilities *without* having to measure the actual value of the spins (which would disturb phase information).

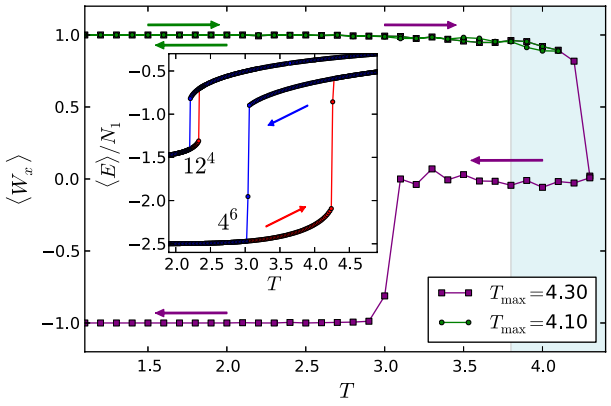

FIG. 2 (color online). Wilson loop in x direction for $L = 4$, $\mathrm{D} = 6$. Each line is a *single* Monte Carlo run, which was warmed from $T = 0$ to $T_{\max}$, then recooled to $T = 0$. Each data point is the average over two thousand Monte Carlo steps at each temperature. Averaging over many Monte Carlo runs would give a purple curve returning to 0 at $T = 0$. The region above $T_c = 3.80$ is shaded. Inset: internal energy for $\mathrm{D} = 4$ and $\mathrm{D} = 6$. For $L = 4$, $\mathrm{D} = 6$, the hysteresis jump occurs at $T = 4.2$.

Mean-field phase diagram for (p, q) codes.—We use mean-field theory (MFT) for Hamiltonian H_A to understand large D behavior. In contrast to Ref. [11], no results will be proven on T_c, but the MFT is still likely exact at large D. From here on, since we consider a Hamiltonian which involves only S^z operators, all our calculations are classical, considering only operators diagonal in the S^z basis, and we set $J_A = 1$. The natural starting point for MFT is a factorized probability distribution, $P(\{S_i^z\}) = \prod_i p_i(S_i^z)$, giving the probability for a spin configuration as a product of the probabilities for each spin, as used in Ref. [14] for the Z_2 gauge theory.

However, this method does not respect gauge invariance. While this invariance only slightly changes T_c for the $(1, \mathrm{D}-1)$ code, it gives a large change in T_c for $(\mathrm{D}/2, \mathrm{D}/2)$ codes where the gauge group is bigger. Gauge invariance is the property that flipping all spins on p-cells attached to any given $p-1$-cell does not change the energy. This leads to an extensive ground state entropy, while there is no way to obtain an extensive entropy in a zero temperature mean-field state. By counting dimensions in a chain complex, the ground state degeneracy of H_A is $2^{N_{p-1}-N_{p-2}+N_{p-3}-\cdots\pm N_0}$ for a (p, q) code, up to $O(1)$ corrections due to any nontrivial homology of the system.

To make a gauge invariant MFT, we add constraints to H_A to fix a unique representative for each gauge group orbit. For a $(1, \mathrm{D}-1)$ code, we use the gauge invariance to fix $S^z = +1$ for all 1-cells oriented along some given lattice direction, which we call the x direction. We use a product ansatz for the remaining spins. We consider only infinite systems here; for a finite system with linear size L with nontrivial homology, we can fix all but a $1/L$ fraction of the spins. Then the Hamiltonian has both four-spin interactions (on plaquettes on which none of the 1-cells are oriented in the x direction) and two-spin interactions (on plaquettes with two 1-cells in the x direction which contain only two unfixed spins). The mean-field equations are

$$\langle S^z \rangle = \tanh[\beta(2\langle S^z \rangle + 2(d-2)\langle S^z \rangle^3)], \tag{3}$$

where $\langle S^z \rangle$ is the average of S^z on the unfixed 1-cells, as can be found by minimizing the variational free energy.

The MFT gives a variational lower bound for the free energy of the system with the spins fixed; subtracting $(1/\beta)N_0 \log(2)$ from this to account for gauge degeneracy gives a lower bound on the free energy $-(1/\beta)\log(Z_A)$. Thus, one could consider several different gauge fixings and choose the one that leads to the lowest free energy. We do not do this here, but it will be useful for the MFT for $(\mathrm{D}/2, \mathrm{D}/2)$ codes where the ground state degeneracy is much larger. For $(1, \mathrm{D}-1)$ codes, the ratio between number of gauge group generators and number of spins goes to 0 as $\mathrm{D} \to \infty$, while for $(\mathrm{D}/2, \mathrm{D}/2)$ codes the ratio approaches $1/2$.

The mean field equations have a trivial solution with $\langle S^z \rangle = 0$ and also have a nontrivial for low T. For large D, the nontrivial mean-field solution exists for $T \leq \mathrm{D}/1.008\ldots$. T_c is strictly smaller than this temperature, as T_c is determined by the crossing of the free energies of the two solutions. The free energy of the trivial solution is $-TN_1 \log(2)$. The free energy of the nontrivial solution requires numerical calculation. For an analytic estimate, we use the free energy of the ground state sector, $-N_2 - TN_0 \log(2) = -[(\mathrm{D}-1)/2 - T\log(2)/\mathrm{D}]N_1$, accounting for gauge degeneracy, giving crossing of the free energies at

$$T_c = D/[2\log(2)]. \tag{4}$$

Corrections to this from numerical solution of the equations are very small as $\langle S^z \rangle$ is close to 1 at the given T_c. The ratio T_c/D in simulations is less than this value, but approaches it as D increases (see Fig. 4).

Monte Carlo measurement of T_c and hysteresis.—Using Metropolis Monte Carlo calculations, we simulate H_A in Eq. (1) on L^D size lattices with periodic boundary conditions. Thermodynamic estimators, e.g., the specific heat, demonstrate that the (1, 2) code has a continuous critical point in the same universality class as the 3D Ising model [15]. For $\mathrm{D} \geq 4$, we observe a strongly first order phase transition, with hysteretic behavior as illustrated in Fig. 2. To obtain an accurate estimate of T_c, one must therefore measure the crossing of two free energy branches, obtained by integrating the internal energy from high and low temperature (see inset of Fig. 3). This procedure accurately reproduces T_c for 3D and 4D (known from duality arguments [14] to be 1.314 and 2.269, respectively) to three decimal places with small system sizes. Figure 4 shows the results for T_c as a function of D, as well as a typical range of hysteresis for $L = 4$.

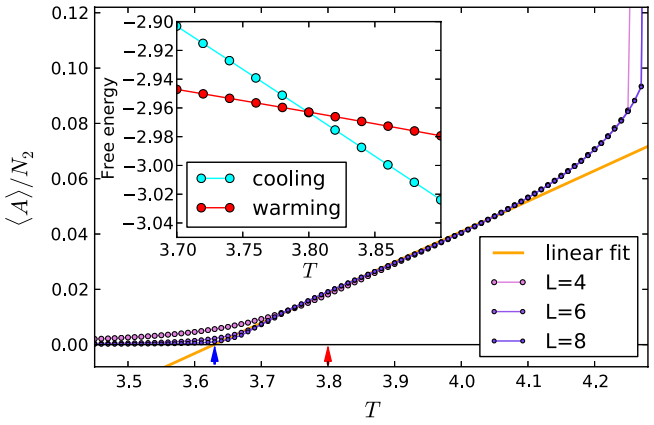

FIG. 3 (color online). Onset of percolation in D = 6. From the crossing of a linear fit to the $L = 8$ data with the axis, we estimate $T_p = 3.63$ (blue arrow on the T axis). Inset: $T_c = 3.80$ (red arrow on the T axis of the main plot) from the crossing of the upper and lower branches of the free energy.

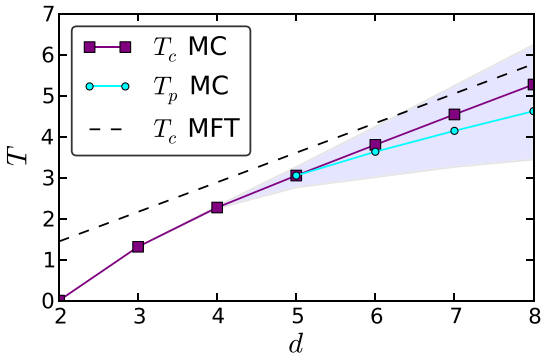

FIG. 4 (color online). Transition temperatures versus dimension. The shaded region represents a typical region of hysteresis in T_c.

We illustrate how one can take advantage of the large hysteresis region to perform error correction even above T_c, by directly measuring the topological bit encoded in a Wilson loop, defined as $W_\alpha = \prod_i^L S_i^z$ for i in a closed line in direction α (where $1 \leq \alpha \leq D$). In Fig. 2, we plot $\langle W_x \rangle$ for two simulations—one where the temperature is increased from $T = 0$ to some $T_{\max} > T_c$ but less than the hysteresis jump, the other for a $T_{\max}$ greater than the hysteresis jump. Upon cooling from $T_{\max}$ to $T = 0$, we observed a near 50% probability that the topological bit was destroyed when $T_{\max}$ was above the hysteresis jump (only one run is illustrated), while for the lower $T_{\max}$, the bit was always retained on the time scale of the simulation. Thus, while we expect that the memory lifetime is *finite* in the thermodynamic limit for $T_c < T < T_{\max}$, the numerical experiments show that information is retained on a very long time scale (much longer than 10^7 Monte Carlo steps per spin at $T = 4.2$ and $L = 4$).

Monte Carlo measurement of percolation.—Monte Carlo simulations are able to give clear estimates of quantities relating to percolation, through measurement of the size of clusters of defect 2-cells (plaquettes with an odd number of $S_z = 1$ on the corresponding 1-cells). Two 2-cells are defined as neighboring if they share the same 3-cell (see Fig. 1). We use this definition to measure quantities related to the size and topology of each defect cluster. In order to identify the unique clusters in a simulation-cell with a given number of defect 2-cells, we developed a variation of the standard Hoshen-Kopelman [16] algorithm for general higher-dimensional networks (see, e.g., Ref. [17]). Using this, we measure the Monte Carlo average of the *largest* cluster size, $\langle A \rangle$, as a function of temperature.

As illustrated in Fig. 3, there is a linear onset of the largest cluster size at some temperature—a clear sign of percolation. To take advantage of hysteresis of the low-temperature phase, simulations were started at $T = 0$ and warmed until the discontinuity in the energy was observed. The percolation transition T_p is determined by extrapolation of the straightest region of the data (Fig. 3). The linear fit matches MFT critical exponents for percolation expected in D ≥ 6. For D ≥ 6, we observe a clear separation of T_p and T_c. For D = 5, we also find a percolation transition in the metastable regime of the low temperature phase, but it appears that $T_p = T_c$ to within simulation errors. For D = 4, no sign of percolation below T_c was seen.

Discussion.—We studied high-dimensional toric codes, finding that self-correction occurs well above T_p, where percolating defects exist. Simulations also find a very large hysteresis region, with the low temperature phase stable on numerically accessible time scales well above the true T_c. In physical implementations, for example, using superconducting circuits [7], active error correction reading error syndromes [18,19] gives nonthermal distributions, but these phenomena may also occur there.

We thank P. Fendley, D. Poulin, and M. Mariantoni for valuable discussions, and especially, L. Hayward and A. Kallin for collaboration on the Monte Carlo software. This research was supported by NSERC of Canada, the Perimeter Institute for Theoretical Physics, and the John Templeton Foundation. Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation. We acknowledge the use of the computing facilities of the Shared Hierarchical Academic Research Computing Network (SHARC- NET:www.sharcnet.ca).

[1] D. Aharonov and M. Ben-Or, *Proceedings of the 29th Annual ACM. Symposium on Theory of Computing* (ACM, New York, NY, 1997), p. 176; E. Knill, R. Laflamme, and W. H. Zurek, Proc. R. Soc. A **454**, 365 (1998); A. Yu. Kitaev, Russ. Math. Surv. **52**, 1191 (1997); D. Gottesman,

Ph.D. thesis, California Institute of Technology [arXiv:quant-ph/9705052].
[2] A. Y. Kitaev, Ann. Phys. (Amsterdam) **303**, 2 (2003); M. Freedman, M. Larsen, and Z. Wang, arxiv:quant-ph/0001108.
[3] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, J. Math. Phys. (N.Y.) **43**, 4452 (2002).
[4] D. Bacon, Phys. Rev. A **73**, 012340 (2006).
[5] J. Haah, Phys. Rev. A **83**, 042330 (2011).
[6] S. Bravyi and J. Haah, Phys. Rev. Lett. **111**, 200501 (2013).
[7] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Phys. Rev. A **86**, 032324 (2012).
[8] See, for example, the calculation of section XC of Ref. [3].
[9] Quantum information can be stored up to the maximum of the two phase transition temperatures: for S^z or for S^x errors. Which of these two temperatures is larger depends on the ratio J_A/J_B in a $(2, d-2)$ code [or more generally, a (p, q) code with $p, q > 1$].
[10] J. L. Lebowitz, A. Mazel, J. Stat. Phys. **90**, 1051 (1998).
[11] J. Bricmont, H. Kesten, J. L. Lebowitz, and R. H. Schonmann, Commun. Math. Phys. **122**, 597 (1989).
[12] V. S. Dotsenko, G. Harris, E. Marinari, E. Martinec, M. Picco, and P. Windey, Nucl. Phys. B **448**, 577 (1995).
[13] R. Peierls, Math. Proc. Cambridge Philos. Soc. **32**, 477 (1936).
[14] A. Wipf, *Statistical Approach to Quantum Field Theory*, Lecture Notes in Physics, Vol. 864 (Springer, 2013).
[15] J. B. Kogut, Rev. Mod. Phys. **51**, 659 (1979).
[16] J. Hoshen and R. Kopelman, Phys. Rev. B **14**, 3438 (1976).
[17] A. Al-Futaisi and T. W. Patzek, Physica (Amsterdam) **321A**, 665 (2003).
[18] A. G. Fowler, D. S. Wang, and L. C. L. Hollenberg, Quantum Inf. Comput. **11**, 8 (2011).
[19] G. Duclos-Cianci and D. Poulin, Phys. Rev. Lett. **104**, 050504 (2010).