Quantum memories at finite temperature

Benjamin J. Brown

*Quantum Optics and Laser Science, Blackett Laboratory, Imperial College London,
Prince Consort Road, London, SW7 2AZ, United Kingdom,
and Niels Bohr International Academy, Niels Bohr Institute,
Blegdamsvej 17, 2100 Copenhagen, Denmark*

Daniel Loss

*Department of Physics, University of Basel,
Klingelbergstrasse 82, CH-4056 Basel, Switzerland*

Jiannis K. Pachos

School of Physics and Astronomy, University of Leeds, Leeds LS2 9JT, United Kingdom

Chris N. Self

*Quantum Optics and Laser Science, Blackett Laboratory, Imperial College London,
Prince Consort Road, London, SW7 2AZ, United Kingdom,
and School of Physics and Astronomy, University of Leeds,
Leeds LS2 9JT, United Kingdom*

James R. Wootton

*Department of Physics, University of Basel,
Klingelbergstrasse 82, CH-4056 Basel, Switzerland*

To use quantum systems for technological applications one first needs to preserve their coherence for macroscopic time scales, even at finite temperature. Quantum error correction has made it possible to actively correct errors that affect a quantum memory. An attractive scenario is the construction of passive storage of quantum information with minimal active support. Indeed, passive protection is the basis of robust and scalable classical technology, physically realized in the form of the transistor and the ferromagnetic hard disk. The discovery of an analogous quantum system is a challenging open problem, plagued with a variety of no-go theorems. Several approaches have been devised to overcome these theorems by taking advantage of their loopholes. The state-of-the-art developments in this field are reviewed in an informative and pedagogical way. The main principles of self-correcting quantum memories are given and several milestone examples from the literature of two-, three- and higher-dimensional quantum memories are analyzed.

CONTENTS

I. INTRODUCTION

Quantum mechanics holds the potential for performing computational simulations (Feynman, 1982) and information processing tasks (Deutsch, 1985; Shor, 1997) much faster than classical technologies. To outperform modern classical machines, a quantum computer must manipulate hundreds of computational qubits to follow a program of millions of quantum logical operations (Hastings *et al.*, 2014; Poulin *et al.*, 2014; Wecker *et al.*, 2014). To realize such a feat using a real physical system, we must preserve a vast entangled quantum state over a long duration while computations are executed. Recognizing that any device exists in an ambient environment at nonzero temperature, we see that probabilistic errors will continually disrupt experimental efforts to coherently control quantum states. It is widely understood that the problem of decoherence is among the largest obstacles impeding the realization of quantum technologies.

The breakthrough that validated the practical possibility of quantum computation was the discovery of quantum error correction (Shor, 1996; Steane, 1996; Lidar and Brun, 2013). The principle behind quantum error correction is to use a redundancy of physical quantum systems to encode a small number of logical computational qubits. We are then able to realize error-correcting protocols by measuring auxiliary physical systems to identify errors affecting the encoded information, and subsequently correct for them. Provided the incident noise is suitably low, we can achieve robust quantum states for an arbitrarily long time using a suitably large redundancy of ancillary systems (Preskill, 1998). While we have recently seen impressive experimental progress in the direction of realizing quantum error-correcting codes (Reed *et al.*, 2012; Barends *et al.*, 2014; Nigg *et al.*, 2014;

Córcoles *et al.*, 2015; Kelly *et al.*, 2015), such active quantum error-correction schemes present an uphill challenge in preparing the complex entangled states that encode quantum information and in reducing incident laboratory noise to the encoded states below a threshold value.

A novel alternative to active error correction would be the discovery of a *self-correcting quantum memory* (Kitaev, 2003; Bacon, 2006), a physical system that is able to reverse the effects of errors by itself. This could be achieved using a many-body Hamiltonian whose energy landscape suppresses large errors that directly affect encoded quantum information. This suppression can be increased indefinitely by increasing the size of the system, allowing arbitrarily large storage times without the need to manually repair the memory.

Ideally, a self-correcting memory must be robust to all forms of physical noise including both finite-temperature effects and small imperfections on the ideal system. Our main focus here will be the study of many-body quantum systems coupled to a thermal bath. To the best of our knowledge, there are no known quantum systems that can preserve coherent quantum information for arbitrarily long time scales at finite temperature, and as such this will be the main consideration of this review (Bacon, 2006). It is also important to recognize that no system will ever be free from weak perturbations such as, for instance, an external magnetic field. Such imperfections may also affect the ability of a system to preserve quantum information so to this end we discuss in parallel known results on the effects of local perturbations on the considered many-body models.

Self-correction is the principle that lies behind the storage of classical information in magnetic media. Here classical bits of information are encoded in the magnetic orientation of some ferromagnetic material. In such a system, thermal noise can cause individual spins to flip, but they will be reoriented quickly by the macroscopic effect of their neighboring spins. As such, the ferromagnet is robust to a spontaneous change of orientation due to the collective behavior of some Avogadro's number of physical spins.

It is an exciting and fundamental question of nature, and indeed the topic of this review, as to whether we can find macroscopic quantum systems to maintain coherent quantum information while simultaneously equilibrating with its surrounding environment. The discovery of such a system will provide a beautiful solution for one of the largest puzzle pieces required to achieve scalable quantum computation. In addition to the remarkable practical applications, the realization of a self-correcting universal quantum computer is also of significant fundamental interest. A macroscopic system that is capable of simulating arbitrarily complex quantum phenomena would provide a powerful demonstration that quantum mechanics is not restricted to only the microscopically accessible parts of the Universe (Farrow and Vedral, 2014).

Many physical systems have been considered for the storage of qubits, for instance, spin qubits in quantum dots (Loss and DiVincenzo, 1998; Kloeffel and Loss, 2013), the ground space of ions (Harty *et al.*, 2014), superconducting systems (Devoret and Scoelkopf, 2013), or other solid-state devices (Fuchs *et al.*, 2011; Saeedi *et al.*, 2013). For a concise review and comparison of different schemes see Schoelkopf and Girvin (2008). A constant challenge for these schemes is

increasing coherence times using mechanisms such as an energy gap to separate excited states from their ground space. Nevertheless, this time scale will always be finite in nature, placing a limit on the computations that can be performed without some error-correction protocol. Ideally, we want to construct a quantum memory that can store quantum states for times that can be tuned arbitrarily using a variable parameter, such as system size.

Furthermore, mediating interactions that entangle qubits encoded in atomic systems requires a coupling bus, for example, the vibrational modes when considering trapped-ion quantum computation or an optical cavity in the case of neutral atoms. These additional structures are subject to thermal errors and are therefore prone to decoherence while performing computational tasks. Ultimately we seek a system that is able to preserve coherent quantum states for time scales much longer than the time it takes to perform logical operations on the encoded states. This would allow the execution of arbitrarily long quantum algorithms given sufficient quantum resources.

Topologically ordered many-body systems (Wen, 2004) play an important role in the study of self-correction. They possess degenerate ground states that cannot be distinguished by local observables. In this feature lies the appeal of topological models as candidate systems for quantum memories; if quantum information is locally indistinguishable, local noise cannot have irreversible effects. Moreover, properties of topologically ordered models have been shown to be stable against weak local perturbations acting on the ideal model Hamiltonian at zero temperature (Kitaev, 2003; Bravyi, Hastings, and Michalakis, 2010). This is an important feature, as we hope that the robust features of topological phases will still be present under realistic conditions where the system will certainly be subject to small imperfections.

In addition to their locally inaccessible degrees of freedom, topological quantum systems are of further interest due to their amenable features for realizing fault-tolerant quantum computation. This has been the subject of intense study (Nayak *et al.*, 2008; Pachos, 2012) for two-dimensional anyonic systems, where quantum information can be stored in collective states of anyons and processed through their braiding. The topological nature of the models again ensures a degree of protection against local noise as long as the anyons are kept well separated. Models that achieve universal fault-tolerant quantum computation by anyon braiding are well known (Freedman, Larsen, and Wang, 2002; Kitaev, 2003; Brennen and Pachos, 2008).

The study of topological quantum computation has extended far beyond the study of anyon braiding. Fault-tolerant computational operations are also realized by the manipulation of holes (Raussendorf, Harrington, and Goyal, 2006; Bombin and Martin-Delgado, 2009; Wootton, 2012), twist defects (Bombin, 2010a; Barkeshli, Jian, and Qi, 2013; Barkeshli *et al.*, 2014), or by other means (Wootton and Pachos, 2011b). It is also noteworthy that research in the direction of computation using experimentally amenable anyon models (Bravyi, 2006; Zilberberg, Braunecker, and Loss, 2008) that do not support a universal set of topological computational operations has led to schemes to supplement such systems with nontopological operations to complete their

computational gate set (Bravyi and Kitaev, 2005; Wootton, Lahtinen, and Pachos, 2009). Consideration of topologically ordered systems as a basis for quantum memories therefore allows us to draw from this wealth of established knowledge to realize a fault-tolerant computational model.

In spite of many known interesting and attractive models, we are yet to rigorously prove the existence of a low-dimensional passively protected quantum memory. It is the purpose of this review to highlight the challenges involved in finding systems that maintain their quantum character at finite temperatures and to discuss new models that come toward a solution to this problem. The present review is separated into two distinct parts. In the first part we introduce the field and paint a picture that demonstrates the difficulty in discovering a stable memory. We show this by means of explicit introductory examples, as well as discussions of rigorously proved no-go theorems for the finite-temperature stability of large classes of systems. In the second part we discuss new models that come some way toward finite-temperature quantum stability over macroscopic time scales. We offer insight into how such models are discovered, and we assess their favorable features and their drawbacks. In doing so we identify underlying open problems and discover established tools that can be used to approach this actively studied and exciting field.

The present review takes the following structure. In Secs. II and III we introduce a common notation, concepts in quantum error correction, and the analytical and numerical methods for examining finite temperature. We conclude Sec. III with a rigorous set of conditions that we demand of a quantum memory, together with a list of attractive features that would make a model suitable for quantum computation and plausible for experimental realization. In Sec. IV, we review the plethora of no-go theorems established so far with respect to passive error correction. We use this study to chart the landscape of the proposed models. The latter Secs. V, VI, VII, VIII, and IX offer a comprehensive review of current actively studied models that demonstrate favorable properties for self-correction. In Sec. X we conclude with an overview of the current state of the field where we discuss open problems that remain unsolved.

II. LOCAL HAMILTONIANS AND QUANTUM ERROR CORRECTION

The study of quantum memories at finite temperature lies at the intersection of the fields of quantum error correction, condensed-matter physics, and statistical mechanics. We therefore require a unifying language that captures the breadth of physics covered by all of these fields. We find such a language in the *stabilizer formalism*. This formalism, initially introduced as an efficient description of quantum error-correcting codes (Gottesman, 2001), provides a natural way of understanding the Hamiltonian models considered here.

The stabilizer formalism efficiently describes quantum error-correcting codes using a list of commuting Pauli operators, known as *stabilizers*. We can use this operator description from quantum error correction to write down a large class of degenerate Hamiltonians. The Hamiltonians we obtain this way have ground-state subspaces that correspond

to the code space of some quantum code, and its excited states reflect the errors the code suffers.

We remark that the stabilizer formalism by no means describes general many-body Hamiltonians that are capable of robust information storage. Indeed, quantum double models (Kitaev, 2003), string-net models (Levin and Wen, 2005; Walker and Wang, 2012), subsystem codes (Poulin, 2005), Turaev-Viro codes (Koenig, Kuperberg, and Reichardt, 2010), and non-Abelian stabilizer codes (Ni, Buerschaper, and den Nest, 2014) are only a few of the classes of interesting Hamiltonian models that are not represented by the stabilizer formalism. In this review, however, we largely restrict our attention to stabilizer models as they provide an analytically tractable class of Hamiltonians upon which many of the developments in this field have been based.

In Sec. II.A we begin by introducing the class of models that we will mainly be concerned with here, namely, *commuting Pauli-Hamiltonians*. In Sec. II.B we give a comprehensive overview of the stabilizer formalism that enables us to identify error-correcting procedures for the considered Hamiltonians. We review how one might perform error correction on either a quantum code or a commuting Pauli-Hamiltonian in Sec. II.C. We then study an explicit and extensively studied example of a commuting Pauli-Hamiltonian in Sec II.D, namely, *Kitaev's toric code model*. In addition to providing a straightforward example of a commuting Pauli-Hamiltonian, the toric code model also exhibits topological order and gives rise to anyonic quasiparticle excitations. We discuss at length the topological nature of the toric code in Sec. II.E. We conclude this section by discussing the stability of the gap at zero temperature in Sec. II.F, a feature presented naturally by topologically ordered systems and an important feature to consider while searching for stable quantum memories.

A. Commuting Pauli-Hamiltonians

We first define the Pauli group $\mathcal{P}_n = \mathcal{P}_1^{\otimes n}$ acting on n distinct two-level quantum systems that we refer to as qubits. The Pauli group $\mathcal{P}_1$ includes, up to phases, the Pauli matrices

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix},$$
$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \qquad (1)$$

and identity $\mathbb{1}$. We often use indices with elements of $\mathcal{P}_1$ to describe the elements of $\mathcal{P}_n$ that act on single qubits. For instance, we write the operator $U \in \mathcal{P}_1$ that acts on the jth physical qubit where $1 \le j \le n$ using the notation $U_j \in \mathcal{P}_n$. Written explicitly, we have

$$U_j \equiv \underbrace{\mathbb{1} \otimes \mathbb{1} \otimes \cdots \otimes \mathbb{1}}_{j-1} \otimes U \otimes \underbrace{\mathbb{1} \otimes \mathbb{1} \otimes \cdots \otimes \mathbb{1}}_{n-j}. \quad (2)$$

This notation is particularly convenient as we can generate the group $\mathcal{P}_n$ up to phases using only the single-qubit operators X_j and Z_j. We finally remark that all elements of $\mathcal{P}_n$ necessarily

have eigenvalues ± 1, which is seen from the fact that $U^2 = \mathbb{1}$ for all $U \in \mathcal{P}_n$.

Having introduced the Pauli group acting on n qubits, we are able to write down Pauli-Hamiltonians that describe interactions between the qubits of a regular lattice. Consider a D-dimensional lattice of qubits of linear size L, as shown in Fig. 1. The $n \sim L^D$ qubits of the lattice are arranged in a structure that depends on the model we introduce. We write down Hamiltonians of the type

$$H = -\frac{\Delta}{2} \sum_j S_j, \qquad (3)$$

where we sum over a set $\mathcal{I}$ of Hermitian interaction terms $S_j \in \mathcal{I}$ such that $\mathcal{I}$ is a subset of $\mathcal{P}_n$.

We must impose physical constraints on Hamiltonian (3). We demand that the Hamiltonian interactions are local. We therefore constrain all elements of $\mathcal{I}$ to have nontrivial, i.e., nonidentity, support only on qubits that can be contained within a box on the lattice of linear size no greater than r, where r is independent of the lattice size. We show a box of linear size $r = 3$ in Fig. 1. Additionally we must bound the interaction strength of the Hamiltonian. To this end we impose that Δ is a positive constant independent of system size. Similarly, we enforce that each qubit supports only a constant number of interaction terms independent of the size of the system.

In general, Hamiltonians that are the sum of local elements of $\mathcal{P}_n$ are intractable for study. We are able to impose additional restrictions that enable us to find solvable classes of Hamiltonians. We first demand that elements $S_j \in \mathcal{I}$ commute, i.e.,

$$[S_j, S_l] \equiv S_j S_l - S_l S_j = 0, \quad \forall\, j, l. \qquad (4)$$

In addition to this, we consider only frustration-free Hamiltonians. Specifically, for all Hamiltonian ground states $|\psi\rangle$, all elements $S_j \in \mathcal{I}$ satisfy the condition

$$S_j |\psi\rangle = (+1)|\psi\rangle, \qquad (5)$$

where ground states are described by an orthonormal basis $|\psi_\mu\rangle$ whose states are indexed by integers μ such that
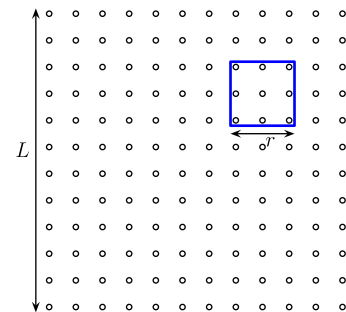

FIG. 1. A regular two-dimensional lattice of qubits with linear size L. All Hamiltonian interactions are contained within a box of linear size r, shown in blue.

$$|\psi\rangle = \sum_\mu c_\mu |\psi_\mu\rangle, \qquad (6)$$

with $\sum_\mu |c_\mu|^2 = 1$. Conditions (4) and (5) enable us to employ the stabilizer formalism that is described in Sec. II.B.

We finally remark on the excited eigenstates of Hamiltonian (3). The Hamiltonian terms $S_j \in \mathcal{I}$ are elements of $\mathcal{P}_n$, and therefore all satisfy the property that $S_j^2 = \mathbb{1}$. It follows from this that the eigenvalues M_j of operators S_j take only values ± 1. We are therefore able to specify excited states of Hamiltonian (3) using the list of eigenvalues M_j. The Hilbert space can then be described completely by the ground eigenspace of H, together with the list of eigenvalues $\{M_j\}$. The excited eigenstates of H are achieved by applying operators $E \in \mathcal{P}_n$ to states $|\psi\rangle$ in the ground eigenspace of H. It is therefore convenient to write Hamiltonian eigenstates $E|\psi\rangle$ and omit any M_j notation. The values M_j are determined by commutation relations $S_j E = M_j E S_j$. The energy eigenvalue ε_E for eigenstate $E|\psi\rangle$ follows immediately from the values M_j such that

$$\varepsilon_E = -\frac{\Delta}{2} \sum_j M_j. \qquad (7)$$

B. The stabilizer formalism

Many quantum error-correcting codes can be described using the stabilizer formalism (Gottesman, 2001). This formalism shares many parallels with the commuting Pauli-Hamiltonians introduced in the previous section. Quantum error-correcting codes describe a subspace of states called the *code space*. We denote an orthonormal basis of states in the code space with vectors $|\psi_\mu\rangle$ for $1 \le \mu \le 2^k$. The code space of a stabilizer code is specified by the *stabilizer group*. The stabilizer group $\mathcal{S}$ is an Abelian subgroup of $\mathcal{P}_n$ where we have defined $\mathcal{P}_n$, the Pauli group for n qubits, in the previous section. A stabilizer group thus defines a quantum error-correcting code such that the code subspace is the common $+1$ eigenspace of all the elements of the stabilizer group. Formally, we write this property such that all elements $S_j \in \mathcal{S}$ satisfy the condition

$$S_j |\psi_\mu\rangle = (+1)|\psi_\mu\rangle, \qquad (8)$$

for all encoded states $|\psi_\mu\rangle$, thus specifying the code space of the code.

The stabilizer group can be described using a generating set of $m \le n$ stabilizers, listed in angle brackets

$$\mathcal{S} = \langle S_1, S_2, \ldots, S_m \rangle, \qquad (9)$$

where all S_j of the generating set are *independent* elements of the stabilizer group, i.e., the stabilizer generators satisfy the condition that $\prod_j S_j^{n_j} = \mathbb{1}$ with $n_j \in \{0, 1\}$ only for $n_j = 0$ for all j. A code of n qubits that is generated by m independent stabilizer generators will encode $k = n - m$ logical qubits.

Encoded logical qubits are manipulated by the group of logical operators $\mathcal{L}$. The group $\mathcal{L}$ is denoted concisely by a generating set of operators $\bar{X}_j$, $\bar{Z}_j \in \mathcal{P}_n$ for $j = 1, 2, \ldots, k$. Operators $\bar{X}_j$ and $\bar{Z}_j$ commute with all elements of $\mathcal{S}$, and with logical operators $\bar{X}_l$ and $\bar{Z}_l$ for $l \neq j$. Operators $\bar{X}_j$ and $\bar{Z}_j$ mutually anticommute. The logical operators therefore generate the Pauli group over the k encoded logical qubits.

We remark that logical operators are not unique with respect to their action upon the code space, but are unique only up to multiplication by stabilizer operators. We consider logical operators $\bar{L}, \bar{L}' \in \mathcal{L}$ that differ only by multiplication by an arbitrary element $S_j \in \mathcal{S}$, i.e., $\bar{L}' = S_j \bar{L}$. Then using the commutation relation $[S_j, \bar{L}] = 0$ we observe that

$$\bar{L}'|\psi_\mu\rangle = S_j \bar{L}|\psi_\mu\rangle = \bar{L} S_j |\psi_\mu\rangle = \bar{L}|\psi_\mu\rangle, \quad \forall\, \mu, \quad (10)$$

thus demonstrating that the actions of $\bar{L}'$ and $\bar{L}$ on the code space are equivalent.

We finally introduce the definition of the weight of an operator and the distance of a quantum error-correcting code. These are useful terms when comparing different error-correcting codes. The weight of operator U, denoted $\mathrm{wt}(U)$, is the number of qubits that U has nontrivial support over. For instance, the operator $U = X_2 X_3$ has $\mathrm{wt}(U) = 2$, as it acts nontrivially on qubits 2 and 3. We use the weight to find the distance d of a quantum error-correcting code. To define the distance we consider least-weight nontrivial, i.e., non-identity, logical operators of a code $\bar{L}^* \in \mathcal{L}$ that satisfy the inequality $\mathrm{wt}(S_j \bar{L}^*) \ge \mathrm{wt}(\bar{L}^*)$ for all elements $S_j \in \mathcal{S}$. The distance of a code is then defined as the weight of the least-weight nontrivial logical operator $\bar{L}^* \in \mathcal{L}$ with the lowest weight. We write this definition concisely such that

$$d = \min_{S_j \in \mathcal{S}} \min_{\bar{L} \in \mathcal{L}} \mathrm{wt}(S_j \bar{L}). \qquad (11)$$

A quantum error-correcting code is able to tolerate and correct for as many as $d/2 - 1$ errors on distinct physical qubits with certainty. In general, however, a code can probabilistically tolerate errors with weight greater than $d/2$, provided the errors incident to the system do not find adversarial configurations with respect to the *error-correction protocol*, as discussed later. The quantum error-correcting codes we review are typically designed to correct low-weight errors with high probability. Here, where correctable errors are discussed, we typically consider errors E such that $\mathrm{wt}(E)/n \ll 1$.

Having introduced the stabilizer formalism, we are now able to explicitly see the correspondence between stabilizer quantum error-correcting codes and frustration-free commuting Pauli-Hamiltonians. The ground states of Hamiltonian (3) are the common $+1$ eigenspace of the set of commuting local interaction terms $\mathcal{I}$. We are therefore able to identify the ground space of Hamiltonian (3) with the code space of a stabilizer group $\mathcal{S}$, whose generators are included in $\mathcal{I}$. In general, $\mathcal{I}$ can be an overcomplete generating set and include some extra elements that are not independent of the others. In Sec. II.A, we specified only that elements of $\mathcal{I}$ are local with respect to the geometry of its underlying lattice.

C. Quantum error-correction protocols

The stabilizers of a quantum error-correcting code are designed to detect the typical errors suffered by encoded quantum states. Provided noise incident to a code occurs at a suitably low rate, we can correctly identify errors with a probability that increases with the distance of the code. This is due to the celebrated *accuracy threshold theorem* (Shor, 1996; Aharonov and Ben-Or, 1997; Kitaev, 1997; Knill, Laflamme, and Zurek, 1998; Preskill, 1998; Aliferis, Gottesman, and Preskill, 2006; Aliferis and Cross, 2007; Terhal, 2015). Once we have identified an error, we can subsequently find an operator that reverses the error and thus corrects for the incident noise. Here we elaborate on the quantum error-correction procedure.

We consider encoded states $|\psi\rangle$ decohering due to a local quantum channel. Given the vast space of realistic noise channels a physical quantum system can suffer, we might suspect that one cannot possibly expect to reverse incident noise. However, if the noise acting on the system is local, and occurs at a sufficiently low rate, then the act of measuring stabilizer operators projects the encoded state onto a state arbitrarily close to $E|\psi\rangle$, where $E \in \mathcal{P}_n$ is some correctable low-weight Pauli error acting on the state. Having measured the stabilizer operators, attempting to determine and correct for the discrete set of Pauli errors E becomes a much more palatable challenge.

In addition to projecting local noise onto an error from the Pauli group $\mathcal{P}_n$, stabilizer measurements S_j also furnish us with information that we can use to estimate the Pauli error E. The set of measurement outcomes $M_j = \pm 1$ for stabilizers S_j are referred to as the error syndrome. Values M_j are determined by the commutation relation

$$S_j E = M_j E S_j, \tag{12}$$

which is seen by consideration of the eigenvalue equation $S_j E |\psi\rangle = M_j E S_j |\psi\rangle = M_j E |\psi\rangle$ that corresponds to the measurement of operator S_j. Obtaining the syndrome data greatly restricts the possible errors that could have occurred, as the incident errors must be consistent with the syndrome information.

There are many errors that can give rise to a given syndrome. To reverse an error E, we consider correction operators $C \in \mathcal{P}_n$ that are consistent with the measured syndrome, i.e., such that $S_j C = M_j C S_j$. If the correction operator satisfies the condition $CE \in \mathcal{S}$, then application of C will restore the quantum error-correcting code to its initial state since $CE|\psi\rangle = |\psi\rangle$ if and only if $CE \in \mathcal{S}$. Alternatively, we may find a correction operator such that CE is a nontrivial logical operator. In this case, we introduce errors that effect the encoded information. We use a *decoder* to attempt to find a correction operator that returns the code to its initial state.

In addition to the error syndrome, the decoder uses information about the error model to find a correction operator that will most likely return the code to its initial state. Specifically, a decoder evaluates the probability $P(\bar{L})$ that the error that caused the syndrome was a member of an equivalence class of errors, where each member is equivalent in the sense that they all have the same effect on the encoded information. Explicitly, the probability that an error is a member of a given equivalence class is determined by

$$P(\bar{L}) = \sum_j \mathrm{prob}(S_j C \bar{L}), \tag{13}$$

where $\mathrm{prob}(E)$ is the probability that Pauli error E is introduced by the known noise model, C is an arbitrary choice of correction operator consistent with the error syndrome, $\bar{L} \in \mathcal{L}$ are the logical operators of the code, and where we sum over all elements $S_j \in \mathcal{S}$. The decoder will then choose the correction operator $C\bar{L}$ as a representative member of the most likely equivalence class to attempt to recover encoded information.

In general it is not always an efficient task to find the most likely equivalence class for which the true error is a member. Instead, we can devise efficient decoding algorithms that approximately determine the most likely class of errors of which the error incident to the code is a member. In Appendix A we describe in detail a specific implementation of an efficient decoder, namely, the clustering decoder, which is commonly used throughout this review. The clustering algorithm is very versatile for decoding the quantum error-correcting codes defined by local commuting Pauli-Hamiltonians.

The correspondence between the syndrome of a quantum error-correcting code and the excited states of commuting Pauli-Hamiltonians means that all the error-correction procedures explained here can be adapted to correct errors suffered by states encoded in the ground space of commuting Pauli-Hamiltonian models.

D. The toric code

We illustrate the concepts introduced using Kitaev's toric code model (Kitaev, 2003). A comprehensive study of the toric code model from the point of quantum error correction can be found in Dennis *et al.* (2002). Qubits are arranged on the edges of a two-dimensional square lattice of linear size L with periodic boundary conditions, i.e., a torus, as shown in Fig. 2. The stabilizer group is generated by star A_v and plaquette B_p, operators for each vertex v, and face p of the lattice. The star and plaquette operators are defined such that

$$A_v = \prod_{\partial j \ni v} X_j, \qquad B_p = \prod_{j \in \partial p} Z_j. \tag{14}$$

Written colloquially, a star operator A_v is the tensor product of Pauli-X operators supported on the edges j that include vertex v in its boundary ∂j, and the plaquette operator is the tensor product of Pauli-Z operators acting on the edges that bound a face p, where the boundary of face p is denoted ∂p. We show examples of such operators in Figs. 2(a) and 2(b). Star and plaquette operators share either zero or two common qubits and therefore commute. The set of all the star and plaquette operators generate the stabilizer group.

When defined on a torus, the toric code model encodes two logical qubits. The logical operators of the model correspond
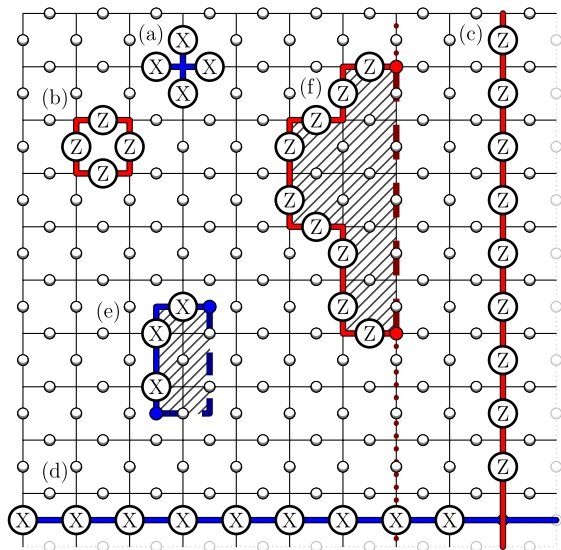
FIG. 2. The toric code lattice. Qubits, shown by white points, are arranged on the edges of a square lattice. The left and dotted right boundaries are unified and similarly the top and dotted bottom edges are unified. (a), (b) A star and plaquette operator, respectively. (c), (d) Logical operators $\bar{Z}_1$ and $\bar{X}_1$, respectively. The Pauli operator in the bottom-right corner of each operator is omitted to show the crossing point. (e) A small error that is easily corrected. The error syndromes are marked by points at the end of the error string. (f) A string error with syndromes separated by half the code distance. We cannot reliably correct this error as there are two available correction operators with equal weight, one of which will lead to a logical error on the code space. The supports of the two distinct lowest-weight correction operators C' and C'' are shown by red dashed and dotted lines, respectively.

to extensive strings of Pauli-X and Pauli-Z operators that wrap around nontrivial cycles of the torus. We show two such operators in Figs. 2(c) and 2(d), respectively. One can see from the diagram that these logical operators have distance $d = L$, the linear size of the system. It is easily checked that these operators commute with all the stabilizers of the code, but mutually anticommute. The displayed logical operators overlap at a single edge of the lattice where, in the diagram, the Pauli operators are omitted.

Error correction for the toric code is particularly intuitive as its syndrome follows a simple geometrical structure. Errors can be regarded as "strings" on the lattice. We show two such errors composed of Pauli-X and Pauli-Z operators in Figs. 2(e) and 2(f), respectively. We cannot detect the positions of the stringlike errors. Instead, the syndrome measurements identify the end points of the stringlike errors. The decoding procedure then consists of using the known end points of the strings and trying to determine the least-weight operator E that may have caused the syndrome. The decoder subsequently returns a stringlike correction operator C that corresponds to a string that reconnects all the stabilizers that returned a -1 measurement outcome. If the errors are few and error strings are very short with respect to the size of the lattice, it is straightforward to identify likely correction operators. This path will connect the end points of string errors such that CE will correspond to a stabilizer operation with high probability,

i.e., $CE \in \mathcal{S}$. In Fig. 2(e) we show a dotted line that supports a suitable correction operator.

In general, the product of an error and its corresponding correction operator will form closed loops on the toric code lattice. The action of these operators will trivially affect the code space only if the loops formed by CE are the boundaries of regions on the lattice. We shade a bounded region in Fig. 2(e).

In the case that either the error strings are very long, such as those shown in Fig. 2(f), or there are many error strings scattered over the lattice, it becomes very difficult to unambiguously find the correction operator C such that no logical error is introduced to the system. For the example given in Fig. 2(f), there are two possible least-weight correction operators of weight $d/2$, which we call C' and C'', whose trajectories are marked by red dashed and dotted lines, respectively. Operator C' is such that $C'E \in \mathcal{S}$. In the diagram we shade the region enclosed by the error and the dashed red line that marks the correction operator. The action of $C''E \in \mathcal{L}$ on the other hand does not enclose a region of the lattice. Instead, as we see, the correction has nontrivial support over an odd number of qubits that support the logical operator shown as a solid blue line at Fig. 2(d). Such a correction will therefore cause a logical error on the code space. It is with this example that we see that determining the correct correction operator becomes difficult once the weight of the error becomes large.

Error correction on the toric code and the structure of its stabilizers can be understood at the fundamental level of homology. This topic goes beyond the scope of the present review, but the interested reader is referred to Nakahara (2003) or Appendix A of Anwar *et al.* (2014) to find a discussion of homology in the context of quantum error correction.

We briefly summarize the quantum error-correcting protocol for the toric code. The system is initialized in a code state by applying appropriate operations; see Fig. 3(a). While the quantum information is stored, errors might occur on the system, as shown in Fig. 3(b). To identify these errors, stabilizer measurements are performed to obtain syndrome information. The locations of the stabilizers that returned a -1 outcome are recorded in Fig. 3(c). A decoder subsequently uses the syndrome information to attempt to find a correction operator. A suitable correction operator that successfully corrects the incident error is applied in Fig. 3(d), thus enabling the reliable readout of encoded quantum information.

E. Topological order and anyons in the toric code

The Hamiltonian of the toric code model (Kitaev, 2003) gives rise to a $\mathbb{Z}_2$ lattice gauge theory (Wegner, 1971; Kogut and Susskind, 1975; Kogut, 1979; Kitaev, 2003; Wen, 2003; Nussinov and Ortiz, 2009a, 2009b). Here we consider the model as a prototypical example of a topologically ordered lattice model with anyonic quasiparticle excitations (Leinaas and Myrheim, 1977; Wilczek, 1982). Its Hamiltonian

$$H_{\text{toric}} = -\frac{1}{2}\sum_v A_v - \frac{1}{2}\sum_p B_p \qquad (15)$$

has degenerate ground states $|\psi_\mu\rangle$ as defined previously. We take interaction strength $1/2$ such that quasiparticles have unit
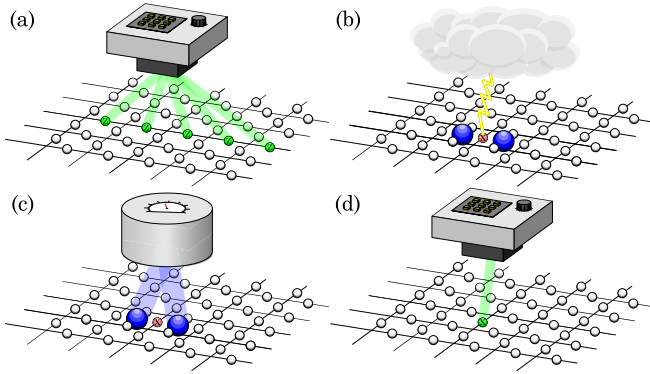
FIG. 3. The error-correcting protocol for the toric code model. (a) The system is initialized. (b) An error occurs due to unavoidable coupling to the environment. (c) The syndrome is measured and fed to the decoding software. (d) The decoding algorithm determines a correction to recover the encoded state and in turn corrects the error.

mass. Its anyonic excitations are a special class of particles that exist in two-dimensional systems. Anyons are of particular interest due to their exchange statistics that are neither fermionic nor bosonic. Interestingly, the concepts of topological quantum error correction and of anyons are intrinsically connected, as will become apparent from the toric code example. For a comprehensive explanation of the general theory of anyons see Appendix E of Kitaev (2006), or alternatively Preskill (2004), Brennen and Pachos (2008), Nayak *et al.* (2008), and Pachos (2012) for an introductory overview. In this section we review the anyonic picture of the excitations of the toric code as it will often provide an efficient description of the dynamics of certain models presented in this review.

The toric code has four types of quasiparticle excitations. The first, known as the vacuum particle, is denoted 1. The vacuum particle describes no anyons. All models, topologically ordered or otherwise, support the vacuum particle. Excited eigenstates of Hamiltonian (15), $|\phi\rangle = E|\psi\rangle$, have electric charges, labeled e, on vertices v that satisfy $A_v|\phi\rangle = -|\phi\rangle$. Similarly, the toric code supports magnetic charges m on faces p whenever $B_p|\phi\rangle = -|\phi\rangle$. The fourth particle of the toric code is known as the dion, labeled ϵ, that is the combination of an e and an m particle.

Anyonic systems have fusion rules to describe the combination of pairs of particles. We write the fusion product of particles a and b as $a \times b$. The fusion product is commutative and associative. For the toric code we have

$$a \times 1 = a,$$
$$e \times m = \epsilon,$$
$$a \times a = 1, \qquad (16)$$

for all $a = 1$, e, m, and ϵ. This anyon model and all others for which the fusion product always leads to a definite result are called Abelian.

In full generality, we can define non-Abelian anyon models, where pairs of anyons can have multiple fusion outcomes.

Like Abelian models, these also require error correction. However, the corresponding error correction problem is quite distinct to that of Abelian anyons, as discussed by Wootton *et al.* (2014). This will have important consequences for the related problem of self-correction. Unlike Abelian anyons, no current proposals for self-correction have been based on non-Abelian models. As such they are beyond the scope of this review. However, recent work on quantum error correction with non-Abelian models can be found in Brell *et al.* (2014), Hutter, Loss, and Wootton (2014), Wootton *et al.* (2014), and Burton, Brell, and Flammia (2015).

Interestingly, quasiparticle excitations of the toric code are created in pairs. We witness this at the microscopic level of the lattice as anyons are created at the two end points of stringlike operators. This feature is reflected by the fusion rule (16), which shows that we require two anyons to recover the vacuum state. Conversely, we can only create anyons from the vacuum in particle-antiparticle pairs.

By using the anyonic description of error operators we find an alternative understanding of the logical operators of the toric code. As described in Sec. II.D, the logical operators are stringlike operators that wrap around nontrivial cycles of the torus. In the anyonic picture stringlike operators correspond to the trajectories of anyons. A logical operator corresponds to the creation of a pair of anyonic particles. One such particle then follows some nontrivial trajectory around the torus and subsequently annihilates with the other pair-created anyon that remained at its initial point. With this point we can define a natural basis for the ground space of the toric code, where orthogonal ground states correspond to different particle fluxes that pass around some arbitrarily selected nontrivial cycle of the torus. We show such a cycle in Fig. 4(a), where the flux of anyon a wraps around the torus along the red line. In the case that many anyonic excitations move around the torus, the ground state is well defined according to the fusion rules of the different particle types. If we change the number of handles, or *genus g*, of the surface where Hamiltonian (15) is embedded, then we change its degeneracy to 2^{2g} and we are able to encode $2g$ qubits there. This is attributed to the extra nontrivial cycles that can be traversed by anyons on the topologically deformed surface.

As an aside remark, the nontrivial braiding statistics between anyons can be obtained from the commutation relations of logical operators (Einarsson, 1990) of the toric code, as the commutation relations between crossing logical

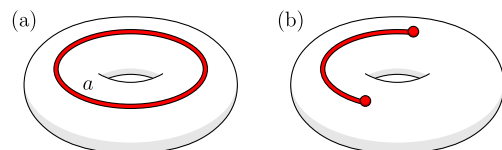FIG. 4. The ground space of the toric code and its low-energy excitations. (a) The ground space of the toric code is naturally described with a basis labeled by anyonic charges a, wrapping around a nontrivial cycle of the torus and then annihilating with its antiparticle, such as that shown in red. (b) Two anyonic excitations created that can propagate at no energy cost to affect the ground space of the system.

operators that follow different nontrivial cycles of the torus correspond to the braiding of anyonic quasiparticles.

Errors can also be interpreted in the anyonic picture. Errors occur when energy is introduced to the system which then creates anyons. Two such anyons are shown in Fig. 4(b). Anyons that propagate around nontrivial cycles on the torus introduce logical errors to the ground space of the system. Unfortunately, once anyons are created on the toric code, it is possible for them to propagate across the system via some suitable mechanism with no additional energy cost. We find this by observing that stringlike operators can be introduced to a ground state of Hamiltonian (15) at constant energy, independent of the length of the string. This insight is the underlying problem that makes it very difficult to design two-dimensional topologically ordered passive quantum memories.

The low-energy excitation of the toric code is an example of a *topological quantum field theory* (Witten, 1988). Models that support topological field theories can be identified by their anyonic statistics, nontrivial ground-state degeneracy, or by means of order parameters such as topological entanglement entropy (Kitaev and Preskill, 2006; Levin and Wen, 2006) that developed from earlier studies of entanglement entropy in topologically ordered lattice models (Hamma, Ionicioiu, and Zanardi, 2005a, 2005b). Topological quantum field theories and extensions thereupon (Hamma, Zanardi, and Wen, 2005; Haah, 2011; Walker and Wang, 2012; Yoshida, 2013) give rise to classes of models that are of interest in the field of quantum memories. In the following section we discuss the stability of the gap that is exhibited by topologically ordered systems at zero temperature in the presence of stray perturbations.

F. Zero-temperature stability

In addition to considering the stability of memories against thermal noise, we must also be mindful of the effects of perturbations when designing quantum memories. Any deviation of a Hamiltonian from our idealized expectations will cause differences in energies as well as dynamics. This can have deleterious effects for any quantum information stored and processed in the system. Fortunately, topologically ordered systems are naturally adept at suppressing the effects of such perturbations at zero temperature. Even so, for arbitrary suppression we require systems of arbitrarily large size, so we are interested in the thermodynamic limit for true stability.

Probably the most well known, but also most misinterpreted result regarding perturbations in topological ordered systems is that of Bravyi, Hastings, and Michalakis (2010). In that work gapped Hamiltonians made up of local, frustration-free, and commuting terms are considered at zero temperature where the degenerate ground-state space is topologically ordered. Local perturbations of general form with finite but sufficiently low strength with respect to the unperturbed Hamiltonian gap are then introduced. It is shown that the splitting of the topologically originated ground-state degeneracy is at most exponentially small with the system size. It is further shown that the gap between the ground-state space and its excited states is also stable against small local

perturbations. The topologically ordered phase is then preserved in the thermodynamic limit, and any given degree of suppression can be efficiently realized. The explicit example of the toric code Hamiltonian perturbed by magnetic fields has been well studied by Trebst *et al.* (2007), Vidal, Dusuel, and Schmidt (2009), Tupitsyn *et al.* (2010), and Dusuel *et al.* (2011).

This result suggests that the time scale at which decoherence by dephasing is induced will diverge exponentially with system size. However, this conclusion is too readily adopted. It is very likely that the system will typically not be in the ground state of the perturbed Hamiltonian. One reason is that, for an arbitrarily large system, it will become a certainty that localized excitations will exist somewhere. Another reason is that the ground state may need to be prepared rather than achieved by cooling. Since the perturbations are not known in general, and since the resultant perturbed ground states may be too complex to prepare, we would expect to use the ground state of the unperturbed Hamiltonian. Finally, perturbations will be time dependent in general.

Since the state of the system will not typically be an eigenstate of the Hamiltonian, the effects of dynamics must be considered. For the case of the toric code, it has been shown that the coherence time will be at most $\mathcal{O}(\log L)$ in the presence of certain local perturbations, including a simple magnetic field (Kay, 2011). Much of this is due to the perturbations enabling anyons to hop across the lattice. It has been shown that this effect can be suppressed by randomizing the couplings of the toric code Hamiltonian, thus introducing Anderson localization (Stark *et al.*, 2011; Wootton and Pachos, 2011a). The lifetime then improves to $\mathcal{O}(\mathrm{poly}(L))$ (Kay, 2011). These dynamical effects have also been studied by Kay (2009), Pastawski *et al.* (2010), Tsomokos, Osborne, and Castelnovo (2011), Bravyi and König (2012), and Röthlisberger *et al.* (2012).

Properties such as these are not necessarily limited to topologically ordered systems. In principle, other types of order may possess equal or perhaps better ground-state stability against unknown perturbations. However, topologically ordered phases are currently the only known means for such suppression and thus form the backbone of current proposals for self-correcting memories.

III. MEMORIES AT FINITE TEMPERATURE

In this section we consider the physics of a quantum memory coupled to a thermal environment. We introduce the necessary mathematical and numerical tools needed to analyze the effects of finite temperature on specific quantum memory models. As a concrete example we analyze the toric code coupled to a finite-temperature environment. We study both qualitatively and quantitatively the time evolution of this system and we identify when the stored information decoheres as a function of bath temperature. We conclude this section with a list of criteria we demand from an experimentally amenable quantum memory.

The present exposition is motivated as a search for systems with quantum properties that are robust at finite temperatures. Nevertheless, the generic thermal dynamics we consider here are widely applicable to other instances of many-body physics

as well as to quantum error correction. For example, it has been understood that quantum error-correcting codes based on self-correcting quantum memories can be decoded locally by using an algorithm based on thermal evolution (Dennis *et al.*, 2002; Pastawski, Clemente, and Cirac, 2011). In a similar spirit, the study of self-correcting memories has also led to the discovery of single-shot error correction (Bombin, 2014). This is a remarkable discovery that could allow us to construct improved quantum error-correcting codes. We discuss local decoders and single-shot error correction in Secs. V.B and IX. C, respectively.

The study of finite-temperature quantum systems is further motivated by the work of Pastawski *et al.* (2009). They showed that an error-correcting code can be passively protected by coupling the code to an auxiliary clock system whose qubits are maintained at infinite temperature. We also point out the work of Kapit, Chalker, and Simon (2014), mentioned in Sec. VI.C. They showed that photon loss at zero temperature in superconducting systems can be modeled as an infinite-temperature noise model in the weak-coupling limit. They draw on this analogy to discover a medium that passively protects quantum information from photon loss in systems where temperature is neglected. From these examples it becomes apparent that the tools and models we develop in this review are broadly applicable to many areas of quantum information and many-body physics.

A. Modeling a finite-temperature environment

Formally, to model a quantum memory at finite temperature, we introduce an auxiliary system, which we call the thermal bath. We couple the bath to the memory system using some appropriate interaction terms that have nontrivial support on both systems. During evolution the interaction terms entangle the memory and the bath. In this way, information stored in the initial state of the memory is shared with the bath and as such it becomes difficult to recover by only accessing the memory.

In general the evolution of a many-body quantum system interacting with a thermal bath is very complicated. In fact it is unknown if the model describing the full thermal evolution is even analytically solvable (Terhal and Burkard, 2005), so to study a memory evolving in a thermal environment we need to make some simplifying assumptions. We assume that the memory interacts weakly with the environment and that the thermal bath is Markovian. A Markovian heat bath is such that the state of the bath is unmodified by interactions with a memory. A consequence of this assumption is that information transferred from a memory to the bath becomes unrecoverable. Additionally, we assume that the bath acts locally on the physical degrees of freedom of the memory. We model the thermal environment such that each qubit is independently coupled to a bath of harmonic oscillators. With this assumption, each event that occurs during the thermal evolution will affect only one physical qubit of the memory system at a time.

In principle, a thermal evolution is accurately described by the bipartite system of the quantum memory and the auxiliary bath. In practice, we need only model the dynamics of the memory subsystem. For this simplification to be valid the time evolution needs to satisfy certain criteria. In particular, the dynamics must evolve the memory toward its Gibbs state

$$\rho_\beta = \sum_j \frac{e^{-\beta \varepsilon_j}}{\mathcal{Z}} |e_j\rangle \langle e_j|, \tag{17}$$

where $\mathcal{Z} = \sum_j \langle e_j | e^{-\beta H} | e_j \rangle$ is the canonical partition function. The vectors $|e_j\rangle$ comprise an orthonormal basis of eigenstates of the memory Hamiltonian, whose corresponding energy eigenvalues are ε_j. We denoted by $\beta = 1/T$ the inverse temperature of the heat bath and we took Boltzmann's constant equal to 1.

An extensive program of research has shown that we can model thermal dynamics of a many-body quantum memory with a simple rate equation (Davies, 1974; Alicki, Fannes, and Horodecki, 2007, 2009; Alicki and Fannes, 2009; Alicki *et al.*, 2010; Chesi, Röthlisberger, and Loss, 2010; Alicki, 2012; Viyuela, Rivas, and Martin-Delgado, 2012; Weiss, 2012). These methods are built from the discovery of exact master equations to study dissipation, a study initially pioneered by Caldeira and Leggett (Calderia and Leggett, 1981; Leggett *et al.*, 1987; DiVincenzo and Loss, 2005). We next summarize the derivation of the dynamical model.

The rate equation evaluates the rate at which an event, described by operator V, occurs during a thermal evolution, such that $\langle e_f | V | e_i \rangle = 1$, where $|e_i\rangle$ and $|e_f\rangle$ are the initial and final eigenstates of the memory with respect to event V. The rate at which event V occurs depends on the difference in energy of the initial and final eigenstates, which we denote as $\omega_V = -(\varepsilon_f - \varepsilon_i)$. We thus have the rate equation that describes the frequency at which event V occurs under thermal evolution

$$\gamma(\omega_V) = \frac{\omega_V}{1 - e^{-\beta \omega_V}}. \tag{18}$$

Intuitively Eq. (18) dictates that processes that increase the energy of the system are exponentially suppressed compared with processes that do not increase the energy of the memory. It is guaranteed that the memory system will evolve toward the Gibbs state if rate Eq. (18) satisfies a detailed balance (Kossakowski *et al.*, 1977). Namely, it must satisfy

$$\gamma(\omega_V) = e^{\beta \omega_V} \gamma(-\omega_V), \tag{19}$$

for all events V. It is easily verified that Eq. (18) satisfies the detailed balance condition.

The open quantum dynamics we have described are derived from a *Lindbladian master equation* (Kossakowski, 1972; Lindblad, 1976). The master equation is obtained by considering the closed dynamics of the system evolving under the Hamiltonian acting on both the memory and the bath subsystems

$$H = H_M \otimes \mathbb{1}_B + \mathbb{1}_M \otimes H_B + \sum_\alpha W_\alpha \otimes f_\alpha. \tag{20}$$

The last term of Hamiltonian (20) describes the interactions between the memory and the bath. Local Hermitian operators W_α and f_α act only on the memory subsystem and the bath subsystem, respectively. Given certain assumptions that we

specify shortly, the evolution of the memory is well described by the master equation

$$\dot{\rho} = i[H_{\mathrm{M}}, \rho] + \mathcal{L}(\rho), \tag{21}$$

where ρ is the density matrix describing the state of the memory subsystem and $\mathcal{L}$ is the Liouvillian. The Liouvillian describes the dynamics due to the interactions between the memory and the bath. It takes the form

$$\mathcal{L}(\rho) = \sum_{\alpha, \omega \geq 0} \mathcal{L}_{\alpha\omega}(\rho), \tag{22}$$

where the individual terms of the Liouvillian are written

$$\begin{aligned}
\mathcal{L}_{\alpha\omega}(\rho) = \hat{g}_{\alpha}(\omega) \{ & V_{\alpha}(\omega)^{\dagger}[\rho, V_{\alpha}(\omega)] + [V_{\alpha}(\omega)^{\dagger}, \rho]V_{\alpha}(\omega) \\
& + e^{-\beta\omega}(V_{\alpha}(\omega)[\rho, V_{\alpha}(\omega)^{\dagger}] + [V_{\alpha}(\omega), \rho]V_{\alpha}(\omega)^{\dagger}) \}.
\end{aligned} \tag{23}$$

In this expression $\hat{g}_{\alpha}(\omega)$ is the power spectrum of the bath and $V_{\alpha}(\omega)$ are the Fourier components of W_{α}, i.e.,

$$U(t)W_{\alpha}U^{\dagger}(t) = \sum_{\omega} V_{\alpha}(\omega)e^{-i\omega t}, \tag{24}$$

where $U(t) = e^{-iH_{\mathrm{M}}t}$.

It can be checked that if we take the interaction terms acting on the memory W_{α} as single-qubit Pauli matrices X_j and Z_j, then the density matrix is diagonal in the energy eigenbasis at any given point in the evolution $\rho(t) = \sum_j p_j(t)|e_j\rangle\langle e_j|$. It can then be shown that Eq. (21) reduces to a much simpler form

$$\dot{p}_k = \sum_{j \neq k} (\Gamma_{j \to k} p_j - \Gamma_{k \to j} p_k), \tag{25}$$

where $\Gamma_{j \to k}$ are the rates of transition from state j to k. Let V be an error process that takes the system from eigenstate j to k with energy cost ω_V. Then the rates are expressed as

$$\Gamma_{j \to k} = \gamma(\omega_V) \propto \frac{\hat{g}(|\omega_V|)}{|1 - e^{-\beta\omega_V}|}. \tag{26}$$

Making the additional assumption that the bath has Ohmic spectral density and large cutoff energy, i.e., $\hat{g}_{\alpha}(\omega_V) \propto \omega_V$ (Leggett *et al.*, 1987; Weiss, 2012) then, up to a normalization, Eq. (26) yields the rate equation (18).

In order to achieve thermalization it is important to require that the interaction operators W_{α} are *ergodic*. This means that the thermal bath is able to address all eigenstates of the memory system. It is known that ergodicity is assured if the only operators that commute with both the memory Hamiltonian and the set of interaction terms W_{α} are proportional to the identity operator (Spohn, 1977; Frigerio, 1978). It is easily checked that if the W_{α} terms are single-qubit Pauli operators then ergodicity is assured.

B. Coherence time of memories

To determine how well a candidate memory performs at finite temperature we need to introduce a suitable figure of merit. To this end we define the coherence time τ as the maximum amount of time information encoded in a system can undergo thermal evolution and remain recoverable with high probability. To evaluate the coherence time of a system, we encode information in a system of interest and evolve the system under the thermal dynamics introduced in the previous section. To recover the information, we allow the use of active error-correction techniques at the time of readout.

To understand the capacity of a system to support quantum information at finite temperature we are primarily interested in the dependence of the coherence time on parameters such as the system size and the inverse temperature of the bath. Naturally, the coherence time will also depend on microscopic details of the system such as the natural units that describe the strength of the local Hamiltonian interactions. These details are overlooked as they are fixed by Eq. (18), but will always take constant values independent of system size.

When evaluating coherence times, we often assume that we can initialize a specific ground state of a system to encode information. This choice is in the interest of providing a fair comparison between different memory systems and also to conceptually simplify our exposition. In general, we expect it to be very hard to prepare a many-body Hamiltonian in its ground state as this will require cooling the system to zero temperature. Alternatively, we might consider manually preparing ground states by means of controlled laboratory operations. However, manual preparation of ground states will also introduce small errors as in general laboratory equipment is fallible (Lodyga *et al.*, 2015). To this end, the ground-state preparation we assume here is unreasonable. However, we do not expect the results we discuss under this assumption to be fundamentally different from the realistic case. Indeed, it is shown in Bombin *et al.* (2013) that random local errors will adjust only the phase transition point of a self-correcting memory.

We also assume that we can realize Hamiltonians that are free from small imperfections such as weak local perturbations as has been discussed in Sec. II.F. Once again, this is not a realistic assumption, as we would typically expect stray fields and other imperfections to alter system Hamiltonians. We make this assumption because the present review is primarily concerned with the finite-temperature behavior of quantum memories. Moreover, this assumption greatly simplifies the computational methods we use to analyze different models. In general, the simultaneous consideration of both temperature and local perturbations makes calculations notoriously difficult, and as such our overview of the field will typically discuss these two forms of noise independently.

We consider now a simple example where we find explicitly the coherence time of a small four-qubit toric code. The four qubits of the model, indexed $j = 1, 2, 3,$ and 4, are subject to the Hamiltonian

$$H_{4\mathrm{Qu.\,toric}} = -\frac{\Delta}{2}(A + B), \tag{27}$$

with stabilizers

$$A = X_1 X_2 X_3 X_4, \qquad B = Z_1 Z_2 Z_3 Z_4. \qquad (28)$$

The code states are given by the four-qubit Greenberger-Horne-Zeilinger states (Greenberger, Horne, and Zeilinger, 1989; Bouwmeester *et al.*, 1999), commonly known as GHZ states, with an even number of qubits in the 1 state. Logical operators for this code act on only two qubits. For example, we can choose $\bar{X} = X_1 X_2$. The thermal error model defined earlier applies single-qubit Pauli operators one at a time. It can therefore apply a logical $\bar{X}$ by first applying X_1 and then X_2. The first operator anticommutes with B and so costs an energy Δ as dictated by Hamiltonian (27). According to the rate equation (18), this process will take a time of around $e^{\beta\Delta}$. The next flip X_2 required to introduce a logical error is a relaxation process and so occurs much more quickly. The coherence time of this small toric code then is $\tau \sim e^{\beta\Delta}$. This is exactly that obtained from Arrhenius' law

$$\tau \sim e^{\beta\varepsilon}. \qquad (29)$$

This law asserts that the coherence time scales exponentially with the energy cost ε of introducing a logical error into the system. In the case of the small toric code, we have that the energy cost of introducing a logical error is equal to the gap of the system $\varepsilon = \Delta$.

Exponential coherence-time scaling with inverse temperature, determined by the constant Hamiltonian interaction strength, is common to all memories of small size. As such, we must look to macroscopic models to find systems with extended coherence times. We must therefore ask what happens to the coherence time of the memory against thermal noise as we increase the system size? It is useful to compare with the benchmark $\tau \sim e^{\beta\Delta}$ for small systems obtained from Arrhenius' law, Eq. (29). The worst possible case for a memory would be sub-Arrhenius scaling of coherence time with β. This would mean that large system sizes have entropic effects that cause the memory to fail faster than for small system sizes. A memory with Arrhenius scaling allows the same protection as one would get against thermal errors for a small system size. Although the resilience of the model to thermal errors may not increase, it may be beneficial to increase the size of the system to improve perturbative stability of a model. We might even expect systems of larger sizes to have greater coherence times than we expect of small system sizes. In which case we might expect super-Arrhenius scaling in coherence time as the temperature is reduced in the limit of large system sizes. Examples of such models are discussed in later sections.

C. The energy barrier

A useful concept in the study of quantum memories is the energy barrier. The four-qubit toric code discussed in the previous section gave an example where Arrhenius's law can be directly applied. In that case we obtained that the lifetime of the memory is correlated in a simple way to its gap Δ. However, the corresponding process for larger many-body systems is much more complicated. Thermal errors act locally,

and so it is not possible to transition between ground states via a single excited state. Instead, errors must navigate a highly degenerate landscape of excited states to modify the ground space of the system. For this reason there is typically no simple generalization to find the value ε that can be used to estimate the lifetime via Arrhenius's law. Nevertheless, we can gain some intuition about the thermalization process by identifying the dominant energy scale of the evolution.

Consider the case in which the commuting Pauli-Hamiltonian system is initially in a logical ground state $|\psi\rangle$. We want to determine how easy it is for thermal errors to rotate the encoded state to, for instance, $\bar{X}|\psi\rangle$, by introducing the logical error $\bar{X}$. As we consider physically motivated local noise models, the logical error operator $\bar{X}$ is decomposed into a sequence of the single-qubit operators U_t that the thermal bath can apply, such that

$$\bar{X} = \prod_{t=1}^{N} U_t. \qquad (30)$$

Here U_t denotes the tth operator to be sequentially applied, and N is the total number of operators required to construct $\bar{X}$. Note that this decomposition of the logical operator into single-qubit operators is not unique. For instance, a permutation in the ordering of the U_t will result in the same action $\bar{X}$ upon state $|\psi\rangle$. Indeed, choosing a different sequence of U_t with different N can yield the same logical operator in the case of commuting Pauli models.

Unlike the initial and final states, the states where the first $0 < t < N$ steps of the error sequence have been applied will be an excited state. Let us use ε_t to denote its energy. For each decomposition of $\bar{X}$ into U_t operators we can consider the energy $\max_t \varepsilon_t$, the maximum energy cost incurred during the sequence. This may be artificially high simply due to a badly chosen sequence. We therefore minimize the energy over all possible sequences to obtain the energy barrier of the model ε_B. This is the minimum energy that the system must achieve in order for a logical error to occur.

We take the toric code Hamiltonian (15) as an example to calculate its energy barrier. A logical operator can be applied by first creating a pair of anyons and then transporting them around a noncontractible loop. The first operation incurs the energy cost for creating a single pair but no subsequent operation will increase the energy of the system, therefore $\max_t \varepsilon_t = 2$. Alternatively, one could generate a logical operator by first applying rotations on every other qubit around a noncontractible loop and then annihilating all the generated anyons by rotating the remaining qubits around the loop. After the creation of all the anyons, the system reaches a state of energy L. We thus have $\max_t \varepsilon_t = L$. Clearly the former logical error path is energetically favorable as it has the smallest energy. Hence, it will be the most common process that introduces logical errors at low temperature. We therefore find the energy barrier of the toric code to be $\varepsilon_B = 2$.

Much of the study of finite-temperature quantum memories has sought complex systems that achieve ε_B that scales with the size of the system. For such systems we should expect their coherence times to scale favorably with system size according to Arrhenius's law. Examples of such models are

studied in Sec. VII. We must keep in mind that in general it is not clear that models can realize logical operators via an ordered sequence of local unitary operators (Haah and Preskill, 2012; Landon-Cardinal and Poulin, 2013). For this reason we reemphasize that the discussion given here is restricted only to commuting Pauli-Hamiltonian models.

D. Free energy and the Curie-Weiss model

In general we cannot completely characterize the coherence-time scaling of a memory by considering only its energy barrier. It is possible that entropic effects can modify the predictions obtained by Arrhenius's law (Temme, 2014; Yoshida, 2014). A more accurate characterization of a system is obtained by consideration of its *free energy*.

With few exceptions, a careful analysis of the free energy of a system is intractable due to its computational complexity. Such an analysis involves careful consideration of an exponentially large number of microstates of the given system. However, evaluating the free energy of a system sheds significant light on its behavior at non-negligible temperatures. In particular, analysis of the free energy enables us to identify low-temperature ordered phases where we expect self-correction to be possible.

The free energy F is the energy cost of an event E, offset by an entropic contribution S, such that

$$F = E - S/\beta. \tag{31}$$

It provides a more accurate estimate of coherence time compared with the energy barrier as it includes the effect of entropy. We obtain the coherence time using the expression

$$\tau \sim e^{\beta F}. \tag{32}$$

In this section we consider a toy model, namely, the Curie-Weiss model, for which the free energy can be evaluated. While unphysical due to its nonlocal Hamiltonian interactions, the Curie-Weiss model is a simple classical model that enables a detailed analysis of the contribution of both the energy barrier and the free energy that can be used to determine its coherence time (Alicki and Horodecki, 2006). For a detailed discussion on the Curie-Weiss model see Kochmański, Paszkiewicz, and Wolski (2013) and references therein.

The Curie-Weiss model is a twofold degenerate model comprised of n classical spins σ_j with $1 \le j \le n$ that take values $\sigma_j = \pm 1$. We denote a configuration of the spins of the system as $\boldsymbol{\sigma}$. The energy of a state of the Curie-Weiss model is

$$E_{CW}(\boldsymbol{\sigma}) = -\frac{\Delta}{n} E_{para}(\boldsymbol{\sigma})^2, \tag{33}$$

where Δ is a constant independent of system size and $E_{para}(\boldsymbol{\sigma}) = -\sum_j \sigma_j$ is the classical Hamiltonian that describes a paramagnet. Specifically, E_{para} assigns one unit of energy to each spin in the -1 state and negates one unit of energy for each spin otherwise. We point out that the nonlocal nature of the Hamiltonian is such that each spin is involved in a number of interaction terms that scales with n. To compensate for this,

the $1/n$ factor in Hamiltonian (33) ensures that the energy cost of a single spin flip does not scale with the size of the system.

The only relevant quantity when studying configurations $\boldsymbol{\sigma}$ is $x = n_\downarrow/n$, where $n_\downarrow$ is the number of spins of configuration $\boldsymbol{\sigma}$ in the -1 state. The two ground states take values $x = 0$ and $x = 1$, and the energy of a typical configuration is

$$E_{CW}(x) = -\Delta n(1 - 2x)^2. \tag{34}$$

It is also important to note that we have $C = n!/(n - n_\downarrow)!n_\downarrow!$ unique configurations that give rise to a particular x. Rearranging, and making use of Sterling's approximation, we obtain

$$C(x) = e^{S(x)}, \tag{35}$$

where

$$S(x) = -n[x \log x + (1 - x) \log(1 - x)] \tag{36}$$

is the entropy of the system. The probability that a system is in configuration $\boldsymbol{\sigma}$ is found using a Boltzmann weight,

$$\text{prob}(\boldsymbol{\sigma}) = e^{-\beta E_{CW}(\boldsymbol{\sigma})}/\mathcal{Z}, \tag{37}$$

where $\mathcal{Z} = \sum_{\boldsymbol{\sigma}} e^{-\beta E_{CW}(\boldsymbol{\sigma})}$ is the partition function of the system. We thus find the probability that a system is in a configuration that takes value x

$$\text{prob}(x) = C(x)e^{-\beta E_{CW}(x)} = e^{-\beta F(x)}/\mathcal{Z}, \tag{38}$$

where now $F(x) = E_{CW}(x) - S(x)/\beta$.

We can use Eq. (38) to understand the behavior of the Curie-Weiss model as a classical memory. We do not require that the memory remains in the ground space to encode a state. We require only that x remains close to its encoded value, where either $x \ll 1/2$ or $x \gg 1/2$. Provided the value of x remains far away from $x \sim 1/2$, we can recover the state of the encoded bit by measuring the magnetization of the system. Finding the magnetization is physically equivalent to taking a majority vote over all the spins of the system. We plot the free energy as a function of x for various β in Fig. 5. The probability that a state takes value x is inversely proportional to the exponent of the free energy, as shown in Eq. (38). Therefore, we can regard the free energy plot of Fig. 5 as a potential landscape, where the system will preferentially find local minima and is unlikely to achieve states with large free energy.

At low temperatures, Fig. 5 shows that the system has two potential minima, one for $x \ll 1/2$, and one at $x \gg 1/2$. At suitably low temperatures, we can increase the depth of the two potential minima by increasing n. As such, it is highly unlikely for a state to achieve a configuration with $x \sim 1/2$ in the thermodynamic limit, as states with a large free energy are achieved very infrequently. Therefore, if we encode a state by preparing it in, for example, a configuration with $x \ll 1/2$, it is unlikely that the thermal environment will evolve the state to one of $x \gg 1/2$ via a sequence of local spin flips, as the evolution must pass through highly improbable states where $x \sim 1/2$. To this end, in the thermodynamic limit, and at
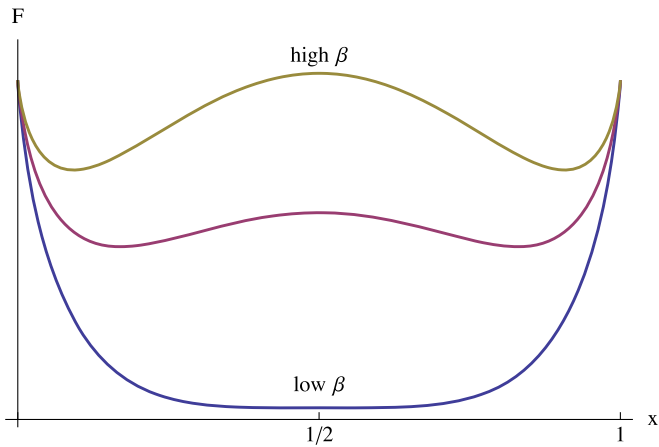
F

FIG. 5. Free energy plotted as a function of x for low, intermediate, and high β, shown by the bottom blue line, the intermediate red line, and the top yellow line, respectively. With decreasing temperature the local minima become more pronounced. At high temperature the entropic contribution is dominant.

suitably low temperatures, the Curie-Weiss model is able to robustly encode a classical bit of information for arbitrarily long time scales. This is shown in Alicki and Horodecki (2006) by taking x to the continuum limit and applying Kramer's formula (Gardiner, 1983).

The Curie-Weiss model is one of the simplest examples of a model that is in an ordered phase at finite temperature. As observed, if the temperature is suitably low, we can robustly encode classical information for arbitrarily long time scales. Conversely, if we increase the temperature, the model undergoes a phase transition into the disordered phase where the storage of information is no longer possible. We observe this in Fig. 5. Specifically, at high temperatures, the free energy curve no longer has two well-resolved minima that are separated by a large potential. Ideally, we seek to find quantum systems that have an ordered phase at finite temperature. We expect that such a system will be able to robustly encode quantum information for long durations. We discuss systems with ordered phases at finite temperature in more depth in Sec. V.

We finally remark that while the free energy offers a much more accurate description of the behavior of a system including evaluation of the phase diagram of a system, it is often difficult to evaluate. As such, we often resort to using simpler concepts such as energy barriers to evaluate the behavior of a system. Recent work has been conducted in this area by Temme (2014), Temme and Kastoryano (2015), and Kómár, Landon-Cardinal, and Temme (2016), where it is shown that Arrhenius's law gives an upper bound on the coherence time of a memory for a large class of local commuting Hamiltonians. Indeed, the resulting coherence times achieved by the use of Arrhenius's law, Eq. (29), is often considered as a widely applicable (Laidler, 1972) rule of thumb.

E. Simulating finite-temperature effects

Monte Carlo methods are frequently used to numerically analyze the evolution of a system where analytical methods

are intractable or to find data that support theoretical conjecture. Here we give an overview of a general method to conduct finite-temperature Monte Carlo simulations for commuting Pauli-Hamiltonians (Bortz, Kalos, and Lebowitz, 1975; Chesi, Röthlisberger, and Loss, 2010).

The noise model approximates the thermal evolution of an eigenstate $|\psi(t)\rangle$ with respect to commuting Pauli-Hamiltonian H when interacting with a thermal bath of inverse temperature β. As mentioned in Sec. II.A, eigenstates of commuting Pauli-Hamiltonians are easily described using a list of eigenvalues M_j. We simulate the noise model as a sequence of discrete events that map between eigenstates of H. At each event we look to obtain some V and δt such that $|\psi(t + \delta t)\rangle = V|\psi(t)\rangle$. At $t = 0$ we typically initialize the state to a ground state of H and we simulate the thermal evolution up to some time $t_{\max}$.

For commuting Pauli-Hamiltonians the random incident errors V are hopping operators that act like X_j, Y_j, or Z_j on the state. Importantly, the action of operators V ensure that $|\psi(t + \delta t)\rangle$ is an eigenstate of H. The relative probability of error event V is determined using Eq. (18). Rates $\gamma(\omega_V)$ are evaluated with respect to H and $|\psi(t)\rangle$. Explicitly, we select error event V by calling from the distribution

$$p_V = \frac{\gamma(\omega_V)}{R}, \tag{39}$$

where we normalize using the total rate $R = \sum_V \gamma(\omega_V)$ with the summation running over all errors V realizable by the noise model.

The time δt that passes between each step as V is applied is determined using R. Since each V occurs as a random process at rate $\gamma(\omega_V)$, the time step δt is a random variable distributed as an exponential distribution with parameter R, the total rate. We numerically generate values of δt such that

$$\delta t = -\frac{\ln(\texttt{rand})}{R}, \tag{40}$$

where `rand` is a random variable chosen uniformly from the interval (0, 1]. We thus obtain the new eigenstate after a time δt, which passes during the event, such that $|\psi(t + \delta t)\rangle = V|\psi(t)\rangle$.

At the end of each event, we check the total time of the system. If $t + \delta t < t_{\max}$ we perform another event using the new eigenstate $|\psi(t + \delta t)\rangle$. Otherwise we stop the simulation and use $|\psi(t + \delta t)\rangle$ and the total incident error to collect sample data.

By averaging over many trials of this process we can obtain estimates of many nonequilibrium thermal quantities, such as the coherence time τ. We use two different methods to estimate τ in this review. In the first we apply the decoder repeatedly as we evolve the system and define τ as the average time it takes for the decoder to fail once. Alternatively we define τ as the time in which the decoder success rate falls below some threshold, e.g., 99%. The values obtained with these different methods may differ by a constant factor, but ultimately will both reveal the coherence-time dependence of the system on its size and temperature.

F. Toric code at finite temperature

The dynamics of the toric code have been extensively studied by Alicki, Fannes, and Horodecki (2009), Freeman *et al.* (2014), Hutter and Loss (2014), and Jouzdani *et al.* (2014). In particular, Alicki, Fannes, and Horodecki (2009) analytically showed using the dynamical model reviewed in Sec. III.A that the coherence time of the model will not exceed $\tau \sim e^{\beta\Delta}$. In this section we provide a self-contained study of the thermal dynamics of the toric code model. We make use of the numerical tools we have discussed throughout this section to complement known analytical results. Importantly, the thermalization dynamics of the toric code provide an explicit example of the behavior of the thermal dynamics of a Hamiltonian system that we aim to defend against.

It is a generic feature of quantum memories that their finite-size behavior differs from their behavior at the thermodynamic limit. Both of these regimes are important. The latter considers properties relevant to the scalability of the system and is pertinent for many theoretical considerations. The former, however, is more relevant to current experimental efforts. Here we use the toric code as a specific example to identify and compare the two different regimes. We identify a critical system size below which finite-size effects are apparent. This size is a function of temperature and the energy cost required to create excitations.

The energy cost to create a single excitation is often referred to as its mass. The mass equates to the interaction strength Δ of the toric code Hamiltonian (15). In thermal equilibrium we expect the average density of anyons in the toric code to scale like $\rho \sim e^{-\beta\Delta}$. Therefore the number of anyon pairs present in a thermalized system of size L is

$$\langle N \rangle \sim \frac{L^2 \rho}{2} = \frac{L^2 e^{-\beta\Delta}}{2}. \qquad (41)$$

Using Eq. (41) we see that for systems smaller than $L \lesssim e^{\beta\Delta/2}$ we have $\langle N \rangle \lesssim 1$. It follows that the probability that there is more than one single pair of anyons on the lattice is negligible. Therefore, in this regime the thermal decoherence of encoded information will most likely occur due to the creation of a single pair of anyons that rapidly propagate across the lattice and introduce a logical error to the memory. We will demonstrate that in the small-size limit the coherence time can be approximated by Arrhenius's law applied to the minimum energy barrier up to system-size dependent corrections, similar to the four-qubit toric code discussed in Sec. III.B. This behavior differs from that of larger system sizes where $L^2 \rho / 2 \sim \langle N \rangle \gg 1$ such that many anyon pairs are uniformly distributed over the lattice. For this case we observe that the coherence time is exponentially shorter than Arrenhius's law predicts and is no longer dependent on system size. The two different limits for the toric code are demonstrated in Fig. 6.

1. Small system size limit

For small systems, and at a temperature low enough for us to expect a good memory, we can typically expect a single pair of excitations to cause the toric code memory to fail. We show
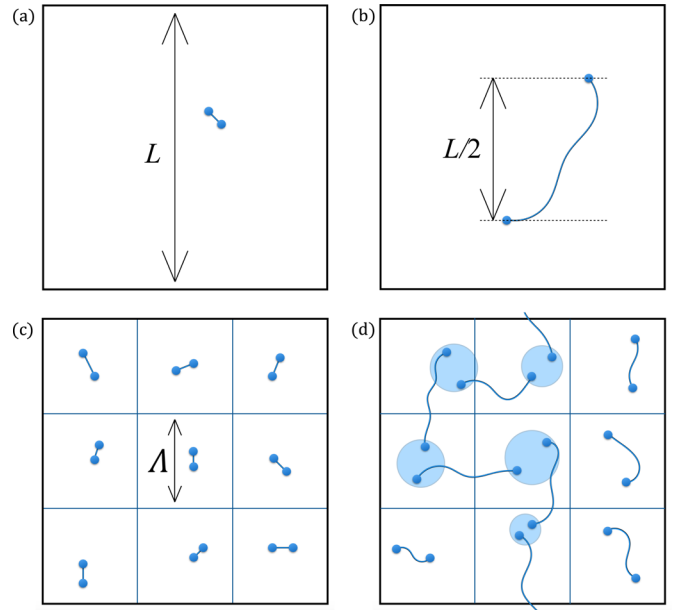

FIG. 6. The thermal dynamics of the excitations of the toric code in both the (a) and (b) small and (c) and (d) large system-size limits. Typical error configurations quickly achieve excitation density $\rho \approx e^{-\beta\Delta}$ for excitations with mass Δ at inverse temperature β. In the small system size limit where $L^2 \rho / 2 \lesssim 1$, it is common for only a single pair of excitations to be created (a), which then rapidly propagate to cause an uncorrectable error (b). In contrast, in the large system-size limit a uniform distribution of anyon pairs is quickly created (c). These pairs diffuse and overlap eventually creating a chain that percolates over the lattice causing an uncorrectable error (d).

such a configuration in Fig. 6(b). We estimate the coherence time $\tau_{\text{small}} = \tau_c + \tau_m$, where τ_c is the time it takes for this pair to be created and τ_m the time it takes for the anyons to diffuse across the lattice up to a separation $L/2$. When the anyons have crossed a distance $L/2$ then the encoded information is irrecoverable by quantum error correction. Only the separation between the anyons is important, so we treat one as fixed and just consider the relative motion. This motion is an unbiased random walk allowing us to easily estimate the diffusion time τ_m. The typical number of steps required for a two-dimensional random walker to reach a distance $L/2$ from its starting point is $(L/2)^2$. Under the noise model discussed in Sec. III.E, the typical time of random walk steps is $[8\gamma(0)]^{-1}$, where the factor of 8 counts the number of processes on the lattice that move anyons. We thus obtain

$$\tau_m \simeq \frac{1}{8\gamma(0)} \left(\frac{L}{2} \right)^2 = \frac{\beta L^2}{32}. \qquad (42)$$

To estimate τ_c we note that the energy cost of pair creation is 2Δ. Again applying the noise model from Sec. III.E, pairs are created from the vacuum at rate $R_0 = 2L^2 \gamma(-2\Delta)$. This implies that the time we wait to see a creation event is $1/R_0 \sim e^{2\beta\Delta}/L^2$. However, not all pairs diffuse to the required distance. Some pairs will instead fuse back to the vacuum at some point later in time; we can quantify the effect this has on

the coherence time by considering the random walks of the pairs.

We denote by $\Pi(L, \beta)$ the probability that a pair does not self-annihilate before reaching separation $L/2$. The motion of the pair is described by an unbiased random walk. It is a standard result for a two-dimensional random walk on a square lattice that the probability of a walker not returning to the origin in the first K steps scales as $1/\ln(K)$. On average we need the walker to avoid self-annihilation for $(L/2)^2$ steps. We therefore expect a factor of $[2\ln(L/2)]^{-1}$ in Π. Additionally, in order to begin the random walk the pair must avoid fusing back to the vacuum immediately. This is a relaxation process that happens at a higher rate than beginning the walk. We therefore include a factor $\sim 1/(1 + A\beta)$ in Π, where A is some constant and $A\beta$ is the relative chance the anyons annihilate when they are nearest neighbors. The probability $\Pi(L, \beta)$ then takes the form

$$\Pi(L, \beta) \sim \frac{1}{1 + A\beta} \frac{1}{\ln(L/2)}. \tag{43}$$

Combining these elements, we expect a creation time scale

$$\tau_c \simeq \frac{1}{R_0} \frac{1}{\Pi} \sim \frac{e^{2\beta\Delta}}{L^2} (1 + A\beta) \ln(L/2). \tag{44}$$

The total coherence time is $\tau_{\text{small}} = \tau_c + \tau_m$. In the small system limit time clearly τ_c is the dominant contribution to τ_{small} due to its exponential dependence on β.

To test these predictions we rigorously study the system evolving in this limit using different numerical experiments with various initial conditions and some variations to the physical noise model. Key technical calculations involved in finding the coherence-time scaling are discussed at length in Appendix B. Here we present the main results concerning the most significant contribution to the coherence time. We separately estimate $\Pi(L, \beta)$, τ_c, and τ_m using numerical simulations. We find good agreement with the predictions of Eqs. (42), (43), and (44) with the constant $A \approx 5$ for the function $\Pi(L, \beta)$. The most significant contribution to τ_{small} comes from τ_c, and the scaling of τ_c is dominated by the factor $1/R_0$. Figure 7 shows the $1/R_0$ scaling of τ_c. This observation matches the predicted values of the key parameters very well, demonstrating a dependence on $1/L^2$ and an exponential growth with $2\beta\Delta$.

The minimum energy barrier of the toric code is 2Δ, giving an Arrhenius law estimate of the coherence time $\tau \sim e^{2\beta\Delta}$. We have shown that in the small-size limit the leading contribution to the coherence time is τ_c, given by Eq. (44). If we ignore the subexponential β dependence inside $1/\Pi$, we can approximate the coherence time by

$$\tau_{\text{small}} \sim e^{2\beta\Delta} \frac{\ln(L/2)}{L^2}. \tag{45}$$

We see that as L becomes larger the lifetime of the toric code memory decreases polynomially up until $L \sim e^{\beta\Delta/2}$. This is a critical size above which the system starts to behave as it would in the thermodynamic limit. We now show that in the
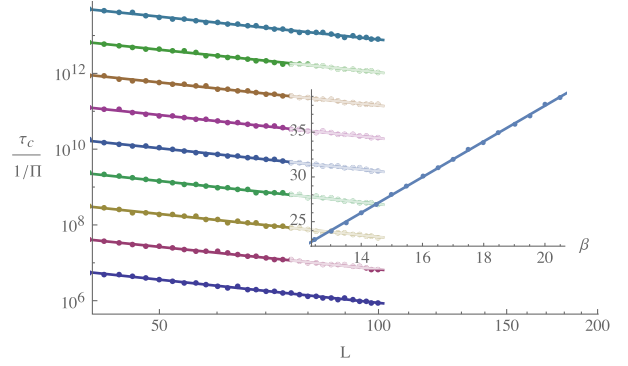

FIG. 7. Pair creation times for toric code excitations in the small system size limit. The average time of pair creation from initialization in the ground state τ_c is shown as a function of system size L, for a range of values from $\beta = 12$ (bottom line) to $\beta = 22$ (top line). The inset shows the values of the fittings in the main plot at the *y*-axis intersection point, as a function of β. Times τ_c are obtained by averaging over 1000 simulations. Here the values of τ_c are divided by a factor of $1/\Pi(L, \beta)$, which we determine numerically independent of τ_c. The gradient of the linear fits, averaged over β, is -2.01, giving an overall scaling of $\tau_c = 0.150(e^{1.99\beta}/L^{2.01})/\Pi$. This is in agreement with the behavior predicted by Eq. (44).

large-size limit the lifetime loses any dependence on system size and is determined only by β.

2. Large system-size limit

In the large-size limit thermalization creates many anyons, with an equilibrium density of $\rho \sim e^{-\beta\Delta}$ for single anyons. On average, anyon pairs are created uniformly throughout the system and each occupies an area of $2/\rho$. We approximate the area as a square of linear size $\Lambda = (\rho/2)^{-1/2} \sim e^{\beta\Delta/2}$ as shown in Fig. 6(c). The probability the decoder fails becomes appreciable once some fraction of pairs separates to distance Λ such that an error chain can percolate through the whole system as shown in Fig. 6(d).

Within a single $\Lambda \times \Lambda$ region the evolution proceeds like the small-size case, i.e., there will be a creation event and the anyons subsequently diffuse apart. We say that a region fails once its anyons move close to anyons from neighboring regions. Assuming each region evolves independently, the time the system decoheres τ_{large} is estimated by the time the typical region fails. This is given by Eqs. (42) and (44) where we set $L/2 = \Lambda$, giving diffusion and creation time scales $\tau_{m,\Lambda} \sim \beta\Lambda^2 \simeq \beta e^{\beta\Delta}$ and

$$\tau_{c,\Lambda} \sim \frac{e^{2\beta\Delta}}{\Lambda^2} (1 + 5\beta) \ln(\Lambda). \tag{46}$$

To get an expression in terms of β we write $\Lambda = Ce^{\beta\Delta/2}$, where C accounts for the constants absorbed into ρ. Then substituting Λ into Eq. (46) and rearranging

$$\tau_{c,\Lambda} \sim e^{\beta\Delta}(1 + r\beta + s\beta^2), \tag{47}$$

where r and s are new constants related to those already introduced. In contrast to the small-size case, the creation and

diffusion time scales have the same exponential dependence on β, and so we expect them both to contribute appreciably to the coherence time. Combining terms we obtain high-temperature coherence time

$$\tau_{\text{large}} \simeq \tau_{c,\Lambda} + \tau_{m,\Lambda} \simeq e^{\beta\Delta}(1 + r'\beta + s\beta^2), \qquad (48)$$

where r' also includes the motional contribution. This simple modeling ignores thermal effects that can now take place such as pair fusion, where two anyons from different pairs fuse to vacuum combining their strings to form a longer string. Processes such as these will alter both the diffusion speed and the pair self-annihilation probability. As a result we do not expect this model to give good predictions of the values of r' and s. However, we do expect it to correctly predict the general features of the dynamics: in particular, the altered exponential dependence on β and the system-size independence of the lifetime. If we ignore its subexponential β dependence the large-size coherence time is

$$\tau_{\text{large}} \sim e^{\beta\Delta}. \qquad (49)$$

This is also exponentially growing with β but at a much slower rate than Arrhenius's law applied to the minimum energy barrier predicts.

We simulate the system evolving in this regime in order to test our assumptions and verify predictions about the dynamics. Here we present the key results. A thorough discussion of our methods and results is given in Appendix B. We verify that anyon densities at the time the decoder fails obey $\rho \sim e^{-\beta\Delta}$ and that for the parameters we consider the typical number of anyons is always large $\langle N \rangle \gg 1$. The average separation between anyon pairs that were either created together or joined by a fusion is seen to scale as $\Lambda \sim e^{\beta\Delta/2}$ as expected. In addition we see that the maximum separation between any pair is always much less than $L/2$ confirming that the decoherence results from the average motion of anyons in local regions. We give numerical data showing the scaling of τ_{large} predicted in Eq. (48) in Fig. 8. Our results are also seen to clearly demonstrate coherence time scaling that is independent of the size of the toric code.

G. Characteristics of self-correcting memories

To properly compare and classify models that are proposed as self-correcting quantum memories, we must have a clear idea of what a self-correcting quantum memory is. We conclude this section by presenting the general characteristics we require from a quantum memory at finite temperature. We use this criteria as a comparative tool to guide us through the presentation of a wide variety of models. We emphasize that the list we give should be regarded as a set of guidelines to be challenged. Other variations of the desiderata asked of a quantum memory are given in Landon-Cardinal *et al.* (2015) and Brell (2016).

In the study of a quantum memory, we first require physically realistic systems. We are interested in Hamiltonian models with interaction terms that are defined locally in three or fewer dimensions. Additionally, we consider only Hamiltonians whose local interaction terms have eigenvalues that are bounded
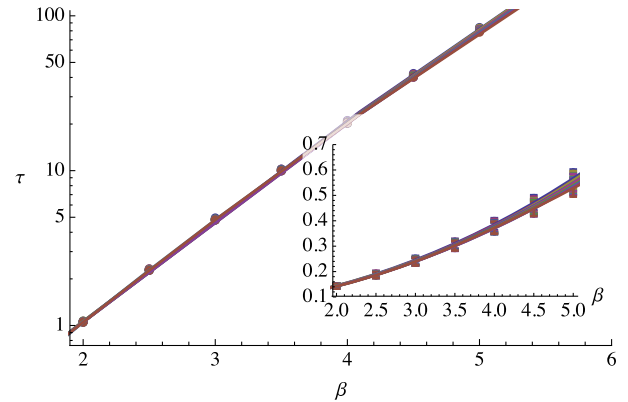

FIG. 8. Coherence time of the toric code in the large system-size limit. Times τ are averaged over 1000 simulations. They are shown here as a function of β for a range of system sizes $L = 100, 120, \ldots, 200$. The data are fit by Eq. (48). We observe only small variations in the fit parameters between system sizes and their averages give the expression $\tau_{\text{high}} = 0.56e^{1.01\beta}(1 + 0.28\beta + 0.31\beta^2)$. The inset shows the nonexponential part of the scaling obtained by dividing the values τ shown in the main plot by $e^{\beta\Delta}$.

by a constant independent of the size of the system. In a similar vein, we also require that each physical degree of freedom in the system supports only a constant number of Hamiltonian interactions. These conditions have been specified for commuting Pauli-Hamiltonians precisely in Sec. II.A.

We next ask what properties we expect of a self-correcting quantum memory. Importantly, we must be able to write information to a quantum memory. We can achieve this using external control during the preparation of the system. Then, encoded information should remain coherent without the application of any control for an arbitrary amount of time while the system is exposed to thermal errors. Ideally we hope that the coherence time of encoded quantum information will diverge to infinity as the size of the system is increased. We require this behavior to be present at some arbitrarily small but nonzero temperature. Such behavior is typically associated with a phase of matter that is ordered at finite temperature below some critical temperature, similar to the ordered phase we observed in Sec. III.D with the Curie-Weiss model.

Further, to ensure that encoded information evolves coherently for an arbitrarily long time, we require that the orthogonal encoded states of a quantum memory are degenerate with respect to the system Hamiltonian. Otherwise, encoded quantum information decoheres rapidly. To this end, we require that the energy splitting between the orthogonal states of the encoded space of the memory vanishes as the size of the system diverges.

In general we do not expect to be able to realize an exact quantum Hamiltonian. Typically a physical system will be subject to minor perturbations due to stray fields or perhaps imperfections in their preparation. We therefore require that the properties that we ask of a quantum memory to be robust against arbitrary local Hamiltonian perturbations, provided the perturbations remain sufficiently weak.

Finally, we require the ability to read out encoded information after the memory has suffered some errors. Even with a

memory with self-correcting properties, we still expect to sustain some small errors that may affect the measurement of logical states. We thus require a decoding algorithm such as those discussed in Sec. II.C and in Appendix A to identify and correct for small physical errors at the point of readout. Moreover, in order for the memory to scale in a practical manner, we require the execution time of the decoding algorithm to scale efficiently with the size of the system.

With these considerations the following list summarizes the criteria we ask of a self-correcting quantum memory:

(1) Locally embeddable in three or fewer dimensions with bounded Hamiltonian interactions and where qubits support a bounded number of interactions.

(2) Encodes a quantum state whose coherence time diverges with system size at a sufficiently low nonzero temperature.

(3) Splitting between the energy levels of the encoded subspace vanishes in the thermodynamic limit.

(4) Memory properties are robust under local perturbations.

(5) Efficiently decodable at readout.

To the best of our knowledge there currently exists no model that has been proven to satisfy all of these criteria. Indeed, we see throughout this exposition that there are many example models that achieve some of these properties, or perhaps weaker notions of these properties, and compromise others. Certainly, it remains interesting to discuss the capability of a quantum memory that does not satisfy all of the listed criteria. We see that some of the models discussed in this review have coherence times that increase with system size up to some cutoff. While not truly self-correcting by our criteria, such models may be very useful if the coherence-time cutoff is large.

While it is difficult to discover a system that satisfies the proposed criteria for a self-correcting quantum memory, it is the goal of this research program to realize a quantum memory in the laboratory that can ultimately be manufactured for the purposes of quantum technologies. Moreover, we require that a memory serves as a component of a larger information processing machine that must work and communicate with other components of a larger processor to complete computational tasks. We therefore append to a list of criteria some additional desiderata that we might reasonably ask of a quantum memory.

We first consider the feasibility of realizing different Hamiltonians. Although it already presents a significant challenge to discover self-correction among Hamiltonians with constant interaction terms, as we phrased the problem, the constant weight of a Hamiltonian interaction can in general be a large constant. In reality, naturally occurring Hamiltonians typically have two-body interactions. It is therefore interesting to discover self-correcting Hamiltonians that have strictly two-body interactions. Similarly, imposing translational invariance is particularly exciting with respect to scalability as we could potentially engineer such a system by designing simple repeating units of the many-body system. We may even expect to find such a system in a regular strongly interacting crystal.

Further, in the interests of experimental amenability, although we can realize three-dimensional systems, such models may be difficult to manipulate. Specifically, we might expect that the quantum degrees of freedom in the center of a three-dimensional crystal will be difficult to access. Such accessibility is likely to be invaluable for encoding and reading out encoded quantum states, and for measuring syndrome data to identify errors suffered by the system. To this end, it is favorable to find a quantum memory in dimensions smaller than 3.

We finally consider fault-tolerant computational abilities in our wish list. Indeed, although finding systems capable of preserving coherent quantum states at finite temperature already presents a considerable challenge, we also want to directly perform interesting computational tasks on information encoded within the memory. Such a property may help reduce computational overhead when we consider manipulating encoded information in a quantum circuit.

We summarize the discussed desiderata as follows:

(1) low-weight, ideally two-body, Hamiltonian interactions;

(2) translational invariance;

(3) embeddable in a low number of dimensions; and

(4) compatible with a fault-tolerant universal quantum gate set.

IV. NO-GO THEOREMS

Before beginning the search for a quantum memory over the vast space of many-body lattice Hamiltonians, it is wise to rule out systems which we cannot expect to maintain quantum information at finite temperature. For this purpose we now consider no-go theorems that identify broad classes of models with physical characteristics that we cannot expect to lead to passively protected memories.

The study of finite-temperature quantum memories requires a breadth of technical aspects, from the abstract mathematical theory of coding to the more physically motivated field of study of finite-temperature effects on lattice Hamiltonians. To this end, no-go results can be broadly separated into two types. We label these general no-go theorems and physically motivated no-go results. The distinction is the following: general no-go theorems seek to exclude large classes of systems from possessing important properties that we expect to be necessary for self-correction. Physically motivated no-go results take into consideration dynamics and microscopic thermal effects to show specific models that will fail to behave well as a quantum memory. Both approaches have complementary advantages and are ultimately of equal importance.

The general no-go theorems typically eliminate the possibility of *energy barriers* in certain classes of systems. Macroscopic energy barriers between degenerate ground states are the basis of our current understanding of finite-temperature stability in classical models. The prototypical case of a classical stable model is the two-dimensional Ising model, which is presented in detail in Sec. V.A. Moreover, it has been shown that an energy barrier is required for a large class of commuting Hamiltonian models (Temme, 2014; Temme and Kastoryano, 2015; Kómár, Landon-Cardinal, and Temme, 2016). It is therefore unlikely that we can expect to find a passive quantum memory with a model that does not support a macroscopic energy barrier.

Physically motivated results lose the generality of their counterpart class of no-go theorems. Instead, they model thermal effects acting on specific models. This approach offers new intuition to show that under physical considerations certain models fail to perform well as a quantum memory. Such results are typically obtained by studying the relevant order parameters that correspond to logical operations acting on the code space of quantum memories. Order parameters are then studied with respect to the dynamics of quantum system when interacting with an auxiliary environmental system, or in its Gibbs thermal equilibrium state. These results support known general no-go theorems for models where it is believed that finite-temperature stability cannot exist.

Known no-go theorems are most limiting in two dimensions. It was shown by Bravyi and Terhal (2009) that two-dimensional commuting Pauli-Hamiltonian models, including the toric code model, cannot support an energy barrier that scales with the size of the system. This result is generalized to two-dimensional topologically ordered commuting Hamiltonians by Landon-Cardinal and Poulin (2013). These results are supported by a wealth of physical no-go results, typically obtained using Kitaev's toric code model, where various order parameters are shown to rapidly decay at finite temperature.

Three-dimensional no-go theorems are significantly less restrictive when compared with their two-dimensional counterparts (Yoshida, 2011; Haah, 2013; Pastawski and Yoshida, 2015). These results leave more promise for the discovery of new models with a macroscopic energy barrier. Indeed, the assumptions necessary to prove the discovered three-dimensional no-go theorems describe a limited set of models when compared with the theorems known in two dimensions. Supporting these results, we also have various physical results showing that topological entanglement entropy (Castelnovo and Chamon, 2008), and the correlation functions of stringlike logical operators (Alicki *et al.*, 2010) decay rapidly for the three-dimensional toric code model. We see in Sec. VII that there are many known physically feasible three-dimensional models that avoid the no-go assumptions we describe here and present favorable properties for finite-temperature stability.

In this section we begin by reviewing no-go theorems in two dimensions. We reproduce the proof of Bravyi and Terhal to show that two-dimensional stabilizer models cannot support a macroscopic energy barrier, and we discuss the supporting physically motivated no-go results. We follow the discussion by considering the no-go theorems in three dimensions. We conclude with possible avenues for avoiding the known no-go theorems. The final section serves as motivation for the positive results that we discuss later on, which include the various models that have been proposed as stable quantum memories.

A. No-go results in two dimensions

In this section we review no-go results in two dimensions. We consider in detail the general no-go theorem due to Bravyi and Terhal (2009), and we discuss the physically motivated no-go results. We motivate this no-go theorem by first considering the toric code, the prototypical model for quantum error correction. We have seen in Sec. II.D that the logical operators of the toric code are one-dimensional stringlike operators. Models with logical operators of this type have a constant energy barrier. To understand this from the anyonic picture of two-dimensional topologically ordered memories given in Sec. II.E, these logical operations correspond to the creation of a pair of anyonic excitations at a constant energy cost which are then free to walk across the lattice at no additional energy penalty. This is discussed in detail by Nussinov and Ortiz (2008) and Alicki, Fannes, and Horodecki (2009), but ultimately follows from the fact that one can find a sequence of single-qubit error operations that will realize a logical operator without increasing the energy of the system beyond a constant value that is independent of the system size. With the toric code in mind it becomes interesting to see if we can find a two-dimensional system with logical operators that are not supported along a one-dimensional line. We follow the proof of Bravyi and Terhal (2009) to show that local two-dimensional stabilizer models necessarily have one-dimensional logical operators which are expected to be incompatible with finite-temperature stability. In the exposition we show how a local noise model can construct a logical operation over its code space at no more than a constant energy cost with respect to the size of the lattice, thus completing the proof.

We consider a two-dimensional square lattice of size $L \times L$. The qubits interact via the local Hamiltonian

$$H = -\sum_j S_j, \qquad (50)$$

where S_j are in general an overcomplete set of local generators for stabilizer group $\mathcal{S} \subset \mathcal{P}_n$ as defined in Sec. II. The ground space of H is the code space of $\mathcal{S}$. Without loss of generality, each local stabilizer generator that acts on a small subset of qubits on the lattice can be contained in a square box no larger than constant linear size r. We show such a box in Fig. 9(a). Moreover, each box can contain no more than a bounded constant number of interaction terms. Hamiltonians with these properties are physically well motivated as described in Sec. II.A. We observe that for the described stabilizer group with local generators S_j, there must exist a logical operator that is supported on a one-dimensional strip of width r, as shown in Fig. 9(a).

We now elaborate on the noise model that introduces a one-dimensional logical error without increasing the energy of the system above a constant independent of its size. The noise model of interest can introduce Pauli errors to single qubits of the lattice. We consider "segments" of a Pauli logical operator supported on the shaded (green) region in Fig. 9(b) which is of variable length $1 \leq l \leq L$, cut off along the horizontal red zigzag line. Importantly, the minimum energy cost associated with this segment is upper bounded by $\varepsilon \sim \mathcal{O}(r^2)$, independent of the system size or l. The part of the logical operator supported on the shaded (green) region will only violate, i.e., anticommute with, stabilizer generators of the physical Hamiltonian that are within a radius $\sim r$ from the cutoff point, shown by the square (blue) area on the lattice. Violated stabilizers correspond to the energy cost of the error on the segment with respect to the Hamiltonian.
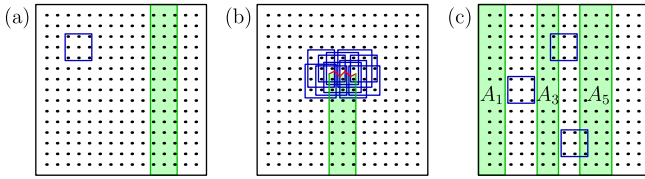
FIG. 9. A sketch of the proof of the no-go theorem due to Bravyi and Terhal. Figures depict square lattices where qubits are marked by black points. (a) Local stabilizer generators are confined to small squares of size r such as that shown in blue in the top-left corner where $r = 3$. For such a code Bravyi and Terhal show that there must exist a logical operator supported on a quasi-one-dimensional strip of width r, such as the vertical strip which is shaded green. (b) An error that forms part of a logical operator, supported on the shaded green vertical strip, does not violate more than a constant number of local stabilizers, marked by blue squares of width r, that extend no further than a distance of $r - 1$ away from the broken end point of the shaded strip, shown as a red zigzag line. (c) A high weight logical operator can be cleaned onto region $A = \bigcup_k A_k$ for odd values of k, where each stabilizer generator, examples of which are supported inside small squares such as those displayed in blue, has common support with no more than one strip. The proof finally uses this fact to conclude that a logical operator must be supported on a single vertical strip.

Given that we have shown that the energy cost of a logical operator segment is upper bounded by energy cost ε independent of l, it suffices to demonstrate that the energy cost of moving the cutoff by a single unit via single-qubit flips for a generic two-dimensional stabilizer model is constant. Indeed, we can change segment length from l to $l + 1$ by overcoming an energy barrier that does not exceed a constant energy cost $\omega \sim \mathcal{O}(r^3)$, independent of system size, before returning to its energy minima ε once the logical segment achieves length $l + 1$. We bound ω by considering the introduction of a single-qubit Pauli error on the lattice. Because of the locality of the terms of the physically constrained Hamiltonian, introducing a new Pauli error can only increase the energy of the system by a constant at most $\sim r^2$. To increase the logical operator segment from length l to length $l + 1$ the noise model introduces a specific set of r single-qubit Pauli errors close to the red zigzag line in Fig. 9(b). This requires the addition of no more than r single-qubit Pauli operators, whose energy cost can be no more than r^2. We are therefore able to bound $\omega \sim r \times r^2$. The described argument holds in the generality of creating a logical operator segment from the ground space of the lattice by considering the increase of the size of a segment from $l = 0$ to $l = 1$.

We have shown that a logical operator segment of length l has energy at most ε, and that we can increase the length of the logical operator segment with energy cost no greater than ω. It follows from this that the single-qubit Pauli error noise model can introduce a logical error to the ground space of the model using a logical operator segment with length $l = L$ with energy never greater than $\varepsilon + \omega$. This demonstrates that commuting Pauli-Hamiltonians with a one-dimensional logical operator have a constant energy barrier.

It only remains to show that a stabilizer group generated locally in two dimensions necessarily has one-dimensional

logical operators. The technical proof makes use of the cleaning lemma, proved by Bravyi and Terhal (2009).

Cleaning lemma. Given a stabilizer group $\mathcal{S}$ that acts on a set of qubits Q, one of the following statements holds for any subset of qubits $A \subseteq Q$:

(1) There exists a nontrivial logical operator $\bar{L} \in \mathcal{P}_n$ supported entirely on A.

(2) All logical operators $\bar{L} \in \mathcal{P}_n$ can be deformed by a stabilizer $S \in \mathcal{S}$, such that $\bar{L}S$ is not supported on A.

We complete the proof using the cleaning lemma. We separate the lattice into an even number of strips of width either r or $r - 1$. We can check that a lattice of size $L = a(r - 1) + br$ can be decomposed into some even number of $a + b$ strips for $L \geq 2(r - 1)^2$.[1] We index the strips in order, and we consider the region of odd strips $A = \cup_{k \in \text{odd}} A_k$, as shown in Fig. 9(c).

We now obtain this proof by contradiction. We assume that there exists a logical operator $\bar{L}' \in \mathcal{P}_n$ whose minimum support cannot be contained on a vertical strip of width r. Because of the width of the strips, such a logical operator cannot be deformed by stabilizers away from region A. Therefore, by the cleaning lemma, it must be possible to find a logical operator $\bar{L} = \bar{L}'S$ for some $S \in \mathcal{S}$ such that $\bar{L}$ is supported entirely on region A. As the logical operator support is wider than a single strip, it must be supported on multiple odd strips A_k. Accordingly, we decompose the logical operator $\bar{L} = \prod_{k \in \text{odd}} \bar{L}_k$, where Pauli operators $\bar{L}_k$ are the support of $\bar{L}$ on strip A_k for odd k.

To complete the argument, we consider operators $\bar{L}_k$. A logical operator will commute with all elements of $\mathcal{S}$. Given the choice of strip width, we observe that the support of any stabilizer overlaps with no more than one odd strip. We show examples of the supports of stabilizer generators within squares in Fig. 9(c). It follows from this fact that, in addition to the logical operator $\bar{L}$, all operators $\bar{L}_k$ must also commute with the stabilizer group. The Pauli operators that commute with the stabilizer group are one of two types of operators. Either they are elements of the stabilizer group, such that $L_k \in \mathcal{S}$, or they themselves are logical operators. Given that $\bar{L}$ is a logical operator, there must be one $\bar{L}_k$ that is a logical operator with width less than or equal to r, providing the desired contradiction. With the observation that we necessarily have at least one logical operator with a one-dimensional support for a stabilizer group which is generated by local two-dimensional stabilizer generators, we conclude the proof that there exists a constant energy barrier between two orthogonal ground states of a commuting Pauli-Hamiltonian in two dimensions.

The discussed work of Bravyi and Terhal was extended in a number of different directions. Landon-Cardinal and Poulin (2013) showed that given a local topologically ordered commuting Hamiltonian there always exists a noise model

---

[1]Proof by induction. We obtain $L = 2(r - 1)^2$ with solution $a = 2(r - 1)$ and $b = 0$. Assume true for $L = a(r - 1) + br$ for all $L \geq 2(r - 1)^2$. In the case that $a > 0$ we obtain $L + 1 = a'(r - 1) + b'r$ with values $a' = a - 1$ and $b' = b + 1$. If $a = 0$ we choose $a' = 2r - 1$ and $b' = b - 2(r - 1) + 1$, which satisfies $a' + b' = \text{even}$, since b is even.

that can locally create a logical operation on the ground space of the model at no more than a constant increase in system energy, thus extending the result of Bravyi and Terhal to a more general class of systems. We remark also on the work of Haah and Preskill (2012) where it is shown that two-dimensional stabilizer models do not support any logical errors that cannot be achieved with an energy cost that is independent of the size of the system under a local noise model.

Supporting the discussed general results in two dimensions, there is also a plethora of physically motivated results in the literature that typically consider the prototypical case: the toric code model. The approaches include a study of topological entanglement entropy (Kitaev and Preskill, 2006; Levin and Wen, 2006) for the toric code model in the thermal equilibrium state; see Castelnovo and Chamon (2007) and Iblisdir *et al.* (2009, 2010). Iblisdir *et al.* (2009, 2010) identified that in realistic systems the topological entanglement entropy vanishes in the large system-size limit at finite temperature for the general class of Kitaev quantum double models (Kitaev, 2003). A further physically motivated study includes the rigorous proof of instability in the toric code model using the Liouvillian open dynamics to show that the expectation values of the logical operators of the toric code model decay rapidly when weakly coupled to a Markovian environment (Alicki, Fannes, and Horodecki, 2009). These results are generalized by Chesi *et al.* (2010) and are considered for the toric code with higher-dimensional spins in Viyuela, Rivas, and Martin-Delagado (2012).

In addition to the physically motivated no-go results, we also remark on the result due to Hastings (2011) which shows that commuting two-dimensional models are unable to support topological order at finite temperature. Similar conclusions were derived by Nussinov and Ortiz (2008, 2009a, 2009b) by consideration of lattice models using methods that were later improved by Chesi *et al.* (2010). Results such as these are particularly important with respect to finite-temperature perturbative stability, which we regard as a required condition for a stable quantum memory. The results of Hastings are supported numerically as discussed by Wootton (2013), where he compares the topological order of unstable memories at finite temperature with stable interacting models. These results are discussed later in Sec. VI.E.

B. No-go results in three dimensions

Thus far we have seen that no-go theorems are very restrictive against energy barriers in two-dimensional systems. It is shown that both commuting topologically ordered Hamiltonians and commuting Pauli-Hamiltonians in two dimensions necessarily support at most a constant energy barrier. In this section we discuss known results in three dimensions. Here the landscape of local Hamiltonians is much more rugged. Indeed, we observed that no-go theorems for three-dimensional models are often superseded by models that avoid the assumptions of a given theorem in a physically sound way.

The first general no-go result in three dimensions is given by Yoshida (2011). Here the methods of Kay and Colbeck (2008) and Yoshida and Chuang (2010) are used to show that three-dimensional Pauli-Hamiltonians that are translationally invariant, and have a constant ground-state degeneracy, must have a one-dimensional logical operator, which can be produced with energy cost independent of the system size.

The work of Yoshida has been supported by physically motivated results that study the three-dimensional toric code model (Hamma, Zanardi, and Wen, 2005) following various approaches. Castelnovo and Chamon (2008) studied the topological contributions to the entanglement entropy of the three-dimensional model in its Gibbs equilibrium state. They found that, while some topological order parameters remain robust up to a critical temperature, a qubit cannot be stored in the three-dimensional toric code at finite temperature because looplike order parameters decay rapidly in the large system-size limit. Similar results are anticipated using the methods of Alicki *et al.* (2010), where they studied the thermal dynamics of the three-dimensional toric code model by considering the model weakly coupled to a Markovian thermal reservoir.

It was explicitly shown that one can surpass the no-go theorem of Yoshida. Haah (2011) constructed a translationally invariant three-dimensional Pauli-Hamiltonian with a macroscopic energy barrier. The model, commonly known as the cubic code, does not have a constant ground-state degeneracy and thus avoids the no-go theorem of Yoshida. We discussed this model in detail in Sec. VII.B.

Following the discovery of the cubic code, Haah generalized Yoshida's no-go theorem (Haah, 2013). He used an elegant representation of the Pauli group to show that translationally invariant three-dimensional commuting Pauli-Hamiltonians can support at best an energy barrier that scales logarithmically with the size of the system. Subsequently, Michnicki (2012, 2014) demonstrated an explicit example of a Pauli-Hamiltonian model that supports a power-law energy barrier, which is not translationally invariant, known as the welded toric code model. In this model the energy barriers of the noncommuting logical operators are varied by changing the size of the system over different length scales. Notably, the welded code has a constant twofold ground-state degeneracy and violates only the translational invariance assumption of the Yoshida proof. We discuss the welded toric code in Sec. VII.C.

Finally, we remark on a recent result given by Pastawski and Yoshida (2015). There they showed that there is a tradeoff for three-dimensional commuting Pauli-Hamiltonians between their capability to support an energy barrier and the fault-tolerant quantum gates that can be achieved by local operations within the ground space of the Hamiltonian. Specifically they considered commuting Pauli-Hamiltonians that can perform a fault-tolerant non-Clifford logical operation by local operations. They show Hamiltonians with such a property cannot support a macroscopic energy barrier. This result is obtained by extending the results of Bravyi and König (2013). An example of such a code which performs a non-Clifford gate, namely, the $\pi/8$ gate, by applying the $\pi/8$ gate locally to each of the physical qubits, is the three-dimensional

color code (Bombin and Martin-Delagado, 2007b). Indeed, this model is not expected to support finite-temperature stability, as it falls into the class of models described by the no-go theorem due to Yoshida.

C. On no-go results

In this section we considered several no-go results. Known no-go theorems identified two large classes of two-dimensional models that cannot support a macroscopic energy barrier; commuting Pauli-Hamiltonians and topologically ordered commuting Hamiltonians. In addition we discussed no-go theorems showing three-dimensional commuting Pauli-Hamiltonians are constrained in their ability to support a power-law energy barrier if they are translationally invariant. These results are summarized in the Venn diagram shown in Fig. 10. Importantly, we provide specific models as examples of the general categories that demonstrate the corresponding behavior.

The no-go results significantly restrict the models we might consider in two dimensions for finite-temperature quantum memories. In particular, it is shown that commuting two-dimensional models cannot support topological order at finite temperature (Hastings, 2011) and conversely that commuting topologically ordered models cannot support a macroscopic energy barrier (Landon-Cardinal and Poulin, 2013). These are very restrictive findings given that we demand perturbative stability for a quantum memory, which is assured with the condition of topological order. Nevertheless, it is not known that topological order is necessary for perturbative stability. As such, there may exist perturbatively stable models that are not topologically ordered that can support a macroscopic energy barrier (Landon-Cardinal and Poulin, 2013). With this in mind, there may still exist commuting two-dimensional models that are suitable as quantum memories.

Another approach to overcoming no-go results in two-dimensional topologically ordered systems is to simply violate their physical assumptions. One study that has attracted notable interest is interacting anyon models. In general, achieving such systems requires the violation of the locality assumption of the discussed no-go theorems. Considerable work has been conducted to find condensed-matter systems that give rise to an effective interacting anyon theory in a local setting. Interacting anyon models are the topic of Sec. VI.

Further, as touched upon in this section, we can obtain positive results for macroscopic energy barriers in three-dimensional commuting Pauli-Hamiltonians. In Sec. VII we review three-dimensional models including the cubic code model, a translationally invariant model with logarithmic energy barrier, and the welded toric code model that obtains a power-law energy barrier by breaking translational invariance.

Curiously, we see in Sec. VII that the logarithmic energy barrier of the cubic code or the power-law energy barrier of the welded toric code do not satisfy the required conditions for a self-correcting quantum memory given in Sec. III.G. To this end, it is even an interesting point of study to try to better understand general necessary and sufficient conditions for models to satisfy the desiderata of a self-correcting memory.
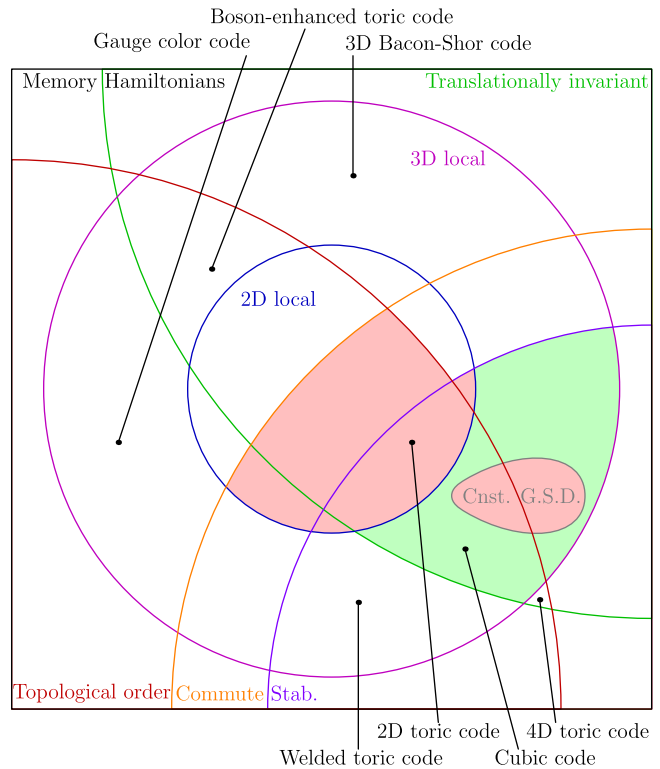

FIG. 10. The landscape of no-go theorems in the space of candidate memory Hamiltonians. Two- and three-dimensional models are shown in magenta and blue circles which are centered in the middle of the diagram. Commuting and stabilizer models are shown inside the orange and purple circles that are centered in the bottom-right corner of the diagram. Translationally invariant models lie within the green circle which is centered at the top-right corner of the diagram. Models satisfying topological order conditions are shown inside the red circle centered at the bottom-left corner of the diagram. Three-dimensional models with a constant ground-state degeneracy lie inside the gray "egg-shaped" region to the right of the figure. Dark red shaded regions have been proven to support an energy barrier that does not scale with the size of the system. Light green shaded areas correspond to models that have energy barriers that scale at best logarithmically with system size. We mark some specific examples of models that we discuss later.

Work in this direction has been conducted by Temme (2014), Temme and Kastoryano (2015), and Kómár, Landon-Cardinal, and Temme (2016).

Beyond commuting Pauli-Hamiltonians there may exist stable quantum memories based on two- or three-dimensional noncommuting Hamiltonians. Because of the difficulty in analytical and numerical calculations for such models, these classes of systems are less well understood compared with their commuting counterparts. However, interesting results have emerged in the field of *subsystem codes*; see, for instance, Bacon (2006), Bravyi and Terhal (2009), and Bravyi (2011). Subsystem codes of particular recent interest with respect to finite-temperature stability include the three-dimensional Bacon-Shor code (Bacon, 2006), the gauge color code (Bombin, 2014, 2015), and the sparse-circuit codes due to Bacon *et al.* (2015). We discuss subsystem codes in Sec. IX.

Finally, we note that the presented general no-go theorems only rule out the possibility of energy barriers in certain classes of models. One might try to sidestep these no-go theorems by finding an alternative method to prevent a finite-temperature environment from corrupting information in the ground space of a quantum memory. We discuss work toward finding such alternatives in Sec. VIII.

In the next section we give consideration to both classical and quantum systems that are known to be thermally stable. Unfortunately, known quantum systems that are proven to be thermally stable are local in dimensions larger than 3. Nevertheless, such analysis gives constructive insight into the properties that give rise to a finite-temperature quantum memory.

V. THERMAL STABILITY IN HIGH DIMENSIONS

Thermal stability in classical systems was first understood by the discovery of the Peierls argument (Peierls, 1936). It shows us that in statistical mechanics stability increases with dimensionality. In this section we follow this trend in the quantum realm. We consider high-dimensional generalizations of well-studied quantum memories to arrive at systems that support finite-temperature stability. Disappointingly, this approach has not yet uncovered a stable model with dimensionality smaller than 4. However, to take a positive outlook on the results summarized in this section, it is demonstrated that finite-temperature stability is not fundamentally inhibited by quantum mechanics. Furthermore, the consideration of high-dimensional quantum models that support finite-temperature stability may enable us to develop a new intuition of thermal stability and may inspire the discovery of a stable quantum memory in lower dimensions.

In this section we review the seminal case of finite-temperature stability by studying the famous two-dimensional classical Ising model. We go on to describe the four-dimensional toric code, the first quantum model rigorously proven to be thermally stable at finite temperature. We conclude this section with a further high-dimensional generalization, namely, the six-dimensional color code model. This model supports both finite-temperature stability and a set of fault-tolerant operations that can implement universal quantum computation.

A. Stability in classical models

In this section, we review the Peierls argument of stability in the two-dimensional classical Ising model (Ising, 1925). The Peierls argument (Peierls, 1936), later refined by Griffiths (1964), shows that the critical phenomena of the Ising model are dependent on its dimensionality. For a modern overview of the Peierls argument, and other important topics relating to the Ising model, we refer the interested reader to Huang (1987) and McCoy and Wu (2014). A modern numerical study of this model is given in the context of a classical memory by Day and Barrett (2012).

While we have more sophisticated methods of extracting the phase diagram of the two-dimensional Ising model due to its exact solution by Onsager (Onsager, 1944; Yeomans, 1992), the intuition developed from Peierls original argument

is a very useful tool for understanding the stability of models where no exact solution is known; see, for instance, Lebowitz and Mazel (1998), Campari and Cassi (2010), and Bonati (2014). Indeed, Peierls argument is used to demonstrate the stability of the high-dimensional quantum systems. It is therefore instructive to give a detailed discussion of Peierls argument applied to the simplest case.

We consider the Ising model defined on an $L \times L$ periodic square lattice of $V = L^2$ spins on its vertices, as shown in Fig. 11. The spin variables σ_j take two values ± 1 and interact via nearest-neighbor interactions described by the classical Hamiltonian

$$E(\boldsymbol{\sigma}) = -\frac{1}{2} \sum_{\langle j,k \rangle} \sigma_j \sigma_k, \qquad (51)$$

where $\langle j, k \rangle$ denote pairs of vertices that are connected by edges of the square lattice, and $\boldsymbol{\sigma}$ is a configuration of all lattice spins. The model behaves as a classical memory that stores a single bit in its twofold degenerate ground space, where the bit is encoded in the magnetization of the system $\bar{\sigma} \equiv \sum_j \sigma_j / V$. The ground states of the model are $\bar{\sigma} = \pm 1$, such that the state $\bar{\sigma} = \pm 1$ corresponds to the configuration where $\sigma_j = \pm 1$ for all j, respectively.

In practice, the Ising model exists at finite temperature. At nonzero temperature the probability of finding the system in the ground state vanishes in the thermodynamic limit. However, for the purpose of storing a bit of information, the stored data will be maintained if the sign of the magnetization is constant. It is expected that the magnetization will maintain the correct sign in the ordered phases of the system. To this end, we must check that the thermal average of the absolute value of the magnetization $\langle |\bar{\sigma}| \rangle$ remains nonzero in the thermodynamic limit for some suitably low but finite temperature.
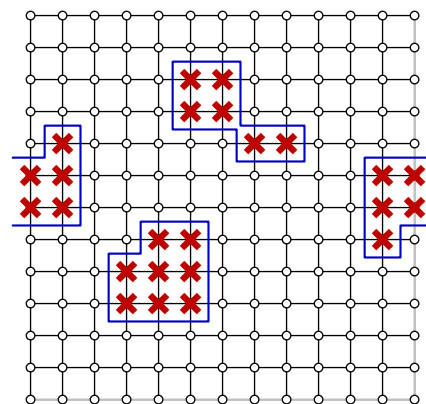

FIG. 11. An example spin configuration of the two-dimensional Ising model. Spins lie on the vertices of a lattice with periodic boundary conditions, and interactions are shown as edges of the lattice. All the spins have the same orientation except for those that have been flipped, which we mark with red crosses. In general, we say that patches of flips occur in "droplets." The energy penalty introduced by a droplet will scale like the length of its boundary. We mark boundaries that enclose droplets with thick solid blue lines which form closed loops.

The argument begins by considering a spin configuration with respect to Hamiltonian (51). In Fig. 11 we have a lattice of spins that are mostly in the $+1$ state, shown by white circles, with regions, or "droplets," of flipped spins in the $\sigma_j = -1$ state. We show three such droplets in the example configuration by patches of red crosses. The Hamiltonian will impose a unit energy cost for each pair of nearest-neighbor spins that have opposite states. As such, the energy cost of a droplet is proportional to the length of its boundary. These boundaries are known as *Peierls contours*, indexed by b, and are marked in blue in Fig. 11. We note that a contour is a single boundary of a droplet, and that in general a given configuration can contain many contours. The probability that a state is in a given configuration from a thermal Gibbs distribution is $p(\boldsymbol{\sigma}) = \exp[-\beta E(\boldsymbol{\sigma})]/\mathcal{Z}$, where $\mathcal{Z} = \sum_{\boldsymbol{\sigma} \in C} e^{-\beta E(\boldsymbol{\sigma})}$ is the partition function and C is the set of all possible configurations. The expectation value of the absolute magnetization is then found by

$$\langle |\bar{\sigma}| \rangle = \sum_{\boldsymbol{\sigma} \in C} |\bar{\sigma}| p(\boldsymbol{\sigma}), \tag{52}$$

which we seek to bound using Peierls argument.

To determine $\langle |\bar{\sigma}| \rangle$, we begin by finding an approximation to the simpler value $\langle N_- \rangle$, the thermal average of the number of spins in the -1 state, for configurations where spins in the $+1$ state are dominant, i.e., states where $\bar{\sigma} > 0$. We refer to regions of $\sigma_j = -1$ as lying "inside" the boundary. In order to evaluate $\langle N_- \rangle$, we first find an upper bound for the number of spins found inside a given contour. Consider a contour of length l. The largest number of flipped spins within such a contour is $l^2/16$. This is because the largest number of flipped spins for a contour of fixed length l occurs when the droplet is square with sides of length $l/4$. For the more general case on the lattice with periodic boundary conditions, one can find droplets that span the lattice with a disjoint boundary of two parts with $l \sim 2L$. In this case, we find an upper bound of $l^2/8$ flipped spins before the number of spins in state $+1$ become dominant. We thus obtain an upper bound for the number of spins in the -1 state for configurations with positive magnetization

$$N_-(\boldsymbol{\sigma}) \leq \sum_b \frac{l_b^2}{8} \delta_b(\boldsymbol{\sigma}), \tag{53}$$

where $\delta_b(\boldsymbol{\sigma}) = 1$ if $\boldsymbol{\sigma}$ contains contour b and 0 otherwise. The term l_b denotes the linear size of contour b.

For Eq. (53) to be meaningful, we must bound the number of contours that have length l. Of course, a boundary must be a closed loop. However, we can find an upper bound for the number of closed loops by calculating the number of random walks of length l that can occur on the lattice, where a walk can begin from any initial position. A walk can begin from one of V possible points. The first step moves in one of four possible directions, and subsequently, to avoid moving backward, we choose from one of three possible directions. Under these conditions we find $4 \times 3^{l-1} V$ possible paths. This method will count each closed loop l times, as a given contour can begin from any of the l faces that the contour crosses. We

therefore arrive at an upper bound for the number of configurations $4 \times 3^{l-1} V/l$. We thus have

$$\langle N_- \rangle \leq \frac{V}{6} \sum_{l=4,6,\dots} l 3^l e^{-\beta l}, \tag{54}$$

where we also used the fact that the thermal average for configurations containing contour b of length $l_b = l$ is suppressed by a Boltzmann factor $\langle \delta_b \rangle \leq e^{-\beta l_b}$. We next take the infinite volume limit to obtain

$$\langle N_- \rangle \lesssim 27 V e^{-4\beta} \frac{2 - 9e^{-2\beta}}{(1 - 9e^{-2\beta})^2}, \tag{55}$$

for $e^{-2\beta} < 1$. By symmetry we find an equivalent value for $\langle N_+ \rangle = \langle N_- \rangle$ over configurations where $\bar{\sigma} < 0$.

We return to the initial problem of obtaining $\langle |\bar{\sigma}| \rangle$. We divide the set of all configurations C into two subsets: C_+ and C_-, where $C_\pm$ contains configurations with a greater number of ± 1 spins. Configurations with $\bar{\sigma} = 0$ will not contribute to magnetization, and we therefore neglect them. We then have that

$$\langle |\bar{\sigma}| \rangle = \sum_{\boldsymbol{\sigma} \in C_+} \bar{\sigma} p(\boldsymbol{\sigma}) - \sum_{\boldsymbol{\sigma} \in C_-} \bar{\sigma} p(\boldsymbol{\sigma}). \tag{56}$$

To complete the argument we use the fact that, by definition, configurations in C_+ have at least $V/2$ spins in the $+1$ state. We can use that $\sum_{\boldsymbol{\sigma} \in C_+} \bar{\sigma} p(\boldsymbol{\sigma}) \geq 1/2 - \langle N_- \rangle/V$. Similarly, we use the relationship $\sum_{\boldsymbol{\sigma} \in C_-} \bar{\sigma} p(\boldsymbol{\sigma}) = \langle N_+ \rangle/V = \langle N_- \rangle/V$ to arrive at

$$\langle |\bar{\sigma}| \rangle \geq 1/2 - 2\frac{\langle N_- \rangle}{V}. \tag{57}$$

We see from Eq. (55) that Eq. (57) has solutions larger than zero for finite values of β, independent of system size, thus demonstrating an ordered phase where the magnetization of the Ising model remains stable in the infinite volume limit of the lattice.

B. High-dimensional stable quantum models

In the previous section we studied suitable conditions for finite-temperature stability by considering the equilibrium state of the two-dimensional classical Ising model. This model is in stark contrast with its one-dimensional counterpart (Ising, 1925), which does not have a finite-temperature phase transition. Instead, it has thermal dynamics akin to those of the two-dimensional toric code model. Indeed, it is a well-understood principle of statistical mechanics that the stability of a model will increase with dimensionality.

Following this reasoning, Dennis *et al.* (2002) showed, using Peierls argument, that the generalized toric code in four dimensions has a finite critical temperature, below which the model is thermally stable. The four-dimensional toric code is defined on a hypercubic lattice. Qubits are placed on the faces of the lattice f. The interactions of the model are six-body operators associated with the links l and the cubes c of the

lattice. Link operators A_l are the tensor product of Pauli-X operators on the faces f which include link l in the boundary of each face ∂f. Similarly, cube operators B_c associated with cube c are the tensor products of Pauli-Z operators on the face qubits that lie on the boundary of the respective cube ∂c. We define the four-dimensional toric code Hamiltonian as

$$H_{4\text{D toric}} = -\sum_l A_l - \sum_c B_c, \qquad (58)$$

where the link and cube operators are given by

$$A_l = \prod_{\partial f \ni l} X_f, \qquad B_c = \prod_{f \in \partial c} Z_f, \qquad (59)$$

respectively. The logical operators of the model are supported in two-dimensional planes which, in four dimensions, intersect at a single point. Each qubit in the model supports four A_l operators, and four B_c operators, such that excitations of the four-dimensional model are not pointlike particles, but instead are linelike particles created by two-dimensional membrane-shaped operators. These geometric features reproduce the energetics of the two-dimensional Ising model that we described earlier in this section. An environment must therefore overcome an $\mathcal{O}(L)$ energy barrier to decohere information encoded in the ground space of the model. Dennis *et al.* (2002) used these features of the model and followed a Peierls argument to show that there is a finite temperature, below which the model lies in an ordered phase.

As an aside, we remark that the argument of Dennis *et al.* (2002) was constructed to show the discussed stable model could be decoded using a local algorithm, i.e., an algorithm that does not require the long-range propagation of classical information. Typically when we consider active error correction, such as the decoder described in Appendix A, we reasonably assume that the classical computation can propagate messages at an infinite velocity when compared with the frequency at which the underlying quantum hardware operates. We are therefore able to design effective decoding algorithms that use syndrome information obtained instantaneously from the entire quantum error-correcting code. Realistically, classical information is communicated at a finite rate bounded by the speed of light. The study of thermally stable quantum memories is therefore interesting from the point of view of decoding algorithms, where benefits might be gleaned from considering quantum error-correcting codes that are analogous to thermally stable memories. The local decoding scheme proposed by Dennis *et al.* has been studied numerically by Pastawski, Clemente, and Cirac (2011). This direction of study has been extended by Harrington (2004) and Herold *et al.* (2014), where a local decoder for the two-dimensional toric code is designed and numerically analyzed. Results in this direction may have important applications from the point of view of local thermally stable quantum memories in low dimensions.

Finally, we remark on extensions to the study of high-dimensional quantum systems. Consideration of Peierls argument suggests that a phase transition occurs at the temperature where the Peierls contours percolate over the system with high probability. The recent work of Hastings, Watson, and Melko (2014) shows, using mean-field arguments and supporting numerical evidence, that as dimensionality increases, the critical temperature of the transition diverges from the temperature at which Peierls contours percolate. This is well understood in the classical case of the D-dimensional Ising model (Lebowitz and Mazel, 1998). The study of high-dimensional quantum memories generalizes known classical results and offers new insights into the physics of phase transitions and critical phenomena.

C. The dynamics of the four-dimensional toric code

Discovering ordered phases as discussed so far in the section only provides a statement about the static equilibrium state of a system. To interrogate the memory time of a quantum memory, one must consider the dynamics of the memory under some realistic evolution.

Alicki *et al.* (2010) rigorously proved that the memory time of the four-dimensional toric code grows exponentially with the size of the system when weakly coupled to a Markovian heat bath. Their results rely on quantum dynamical semigroups (Alicki and Lendi, 2007). The tools that were developed to derive their results were built over a series of papers (Alicki and Fannes, 2009; Alicki, Fannes, and Horodecki, 2009). We point out that the results we summarize in this section are generalized and simplified in Chesi *et al.* (2010) and Bombin *et al.* (2013), where subsystem codes and high-dimensional color codes are, respectively, considered.

The thermal evolution of a many-body quantum state is very difficult to analyze. To simplify the problem, Alicki *et al.* (2010) studied the evolution of an anticommuting pair of observables $\tilde{X}$ and $\tilde{Z}$ that we specify shortly, which act on a two-dimensional subspace of the Hilbert space of the physical system. Specifically, they related the fidelity of the qubit acted upon by the observables $\tilde{X}$ and $\tilde{Z}$ to their decay rates λ, which is defined with respect to some Liouvillian $\mathcal{L}$ as is given in Eq. (17). The decay rate of observable O is defined

$$\lambda(O) = -\text{tr}[\rho_\beta O^\dagger \mathcal{L}(O)], \qquad (60)$$

for observables satisfying $\text{tr}(\rho_\beta O) = 0$ and $\text{tr}(\rho_\beta O^\dagger O) = 1$, where ρ_β is the Gibbs state of the system; see Eq. (17). The interaction terms of the Liouvillian considered by Alicki *et al.* (2010) are single-qubit Pauli operators.

The work of Alicki *et al.* then showed that the fidelity F of an encoded state $\rho(t)$ decays over time like

$$F(\rho(t)) \equiv \langle \psi | \rho(t) | \psi \rangle \geq \tfrac{1}{2}(e^{-\lambda(\tilde{X})t} + e^{-\lambda(\tilde{Z})t}), \qquad (61)$$

with respect to the initial pure state $\rho(0) = |\psi\rangle\langle\psi|$.

This discussion reduces the problem of finding the coherence time of the four-dimensional toric code to finding an upper bound for $\lambda(\tilde{X})$ and $\lambda(\tilde{Z})$. In order to do so, we must first describe the dressed logical operators $\tilde{X}$ and $\tilde{Z}$ for the four-dimensional toric code. A dressed logical operator takes the form

$$\tilde{X} = \bar{X}C^X, \qquad \tilde{Z} = \bar{Z}C^Z, \qquad (62)$$

where $\bar{X}$ and $\bar{Z}$ are the two-dimensional logical operators of the four-dimensional toric code. The operators C^X and C^Z, rigorously defined by Alicki *et al.* (2010), Chesi *et al.* (2010), and Bombin *et al.* (2013), effectively performed the role of error correction, as discussed in Sec. II.C. Specifically, operator C^X first projects the encoded state onto an eigenstate of the B_c operators, defined in Eq. (59), and subsequently applies a *low-weight* Pauli correction operator that returns the system to the $+1$ eigenspace of the B_c operators. Similarly, C^Z projects the system onto an eigenstate of the A_l operators, and subsequently applies a low-weight correction operator that returns the system onto the $+1$ eigenspace of the A_l operators. The low-weight correction operators of the C^X and C^Z operators are obtained efficiently from a configuration of eigenvalues of A_l or B_c operators using the clustering decoder described in Appendix A.

Having defined the dressed observables of the four-dimensional toric code, it remains only to upper bound the decay rates $\lambda(\tilde{X})$ and $\lambda(\tilde{Z})$ to show that the fidelity of an encoded qubit decays slowly if the system is large. Because of the symmetry between $\tilde{X}$ and $\tilde{Z}$, we restrict our attention to only the $\tilde{X}$ operator. An equivalent discussion holds for the $\tilde{Z}$ operator.

To bound $\lambda(\tilde{X})$, an extension of Peierls argument, discussed in Sec. V.A, is employed. First it is shown that the equilibrium state of the four-dimensional toric code at a suitably low temperature is dominantly populated by configurations of small loop excitations that are created by low-weight configurations of errors. Such an equilibrium state can be successfully corrected by the C^X operator with arbitrarily high probability. Indeed, it is easily checked using Peierls argument that the probability of observing loop excitations that are larger than a specified length that is a fraction of the linear size of the system is exponentially suppressed (Dennis *et al.*, 2002).

Finally, to understand the dynamics of the equilibrium state of the four-dimensional toric code one must consider the action of the Liouvillian on the dressed observable. Once again, an extension of Peierls argument is used to show that the probability that local errors will introduce large looplike excitations to the Gibbs state is exponentially suppressed in the size of the system. It follows that the decay rate of both $\tilde{X}$ and $\tilde{Z}$ operators are exponentially suppressed, thus providing the desired result by application of the decay rates to Eq. (61).

This argument demonstrates that the four-dimensional toric code Hamiltonian given in Eq. (58) is self-correcting at a sufficiently low temperature. Moreover, the four-dimensional toric code satisfies the conditions required to demonstrate perturbative stability at zero temperature by the proof given by Bravyi, Hastings, and Michalakis (2010). It will be interesting to show that the four-dimensional toric code, or indeed a quantum system of any dimensionality, is self-correcting at finite temperature, even in the presence of weak local perturbations. Hastings (2011) proposed a definition for topological order at finite temperature and showed that it is satisfied by the four-dimensional toric code. One approach to demonstrating that the self-correcting properties of a model are preserved under local perturbations might be to determine if the definition proposed by Hastings implies that a system is perturbatively stable.

D. Thermally stable quantum computation

We conclude this section with a discussion on the more general problem. What is the smallest dimensionality where we obtain both thermal stability and the desirable feature of a gate set that can be executed fault tolerantly to realize universal quantum computation? This problem has been approached by Bombin *et al.* (2013). They considered D-dimensional generalizations of the color code models (Bombin and Martin-Delgado, 2007a). These models are of particular interest due to the extended set of gates they can achieve on their ground states transversally.

A logical gate on the ground space of the code is executed transversally when one can make a logical rotation on the code space of a code by applying local rotations to its physical degrees of freedom. This is a favorable approach to performing gates as local operations on individual degrees of freedom do not propagate errors during their application. In general, the available transversal gates of a given model are limited by its microscopic details. Notably, the two-dimensional color code (Bombin and Martin-Delgado, 2006) can perform the Clifford gate set transversally. Together with the noisy preparation of magic states (Bravyi and Kitaev, 2005), the Clifford gate set achieves universal quantum computation. It has also been discovered that a three-dimensional color code can achieve fault-tolerant universal quantum computation (Bombin and Martin-Delgado, 2007b). Transversally, this three-dimensional model can achieve a $\pi/8$ gate, and a controlled-not gate. Together with the ability to prepare and measure the ground space in the logical X and the Z basis, this model achieves universal quantum computation. Sadly, however, the three-dimensional color code does not support finite-temperature stability, as shown by the no-go theorem due to Yoshida (2011).

To achieve a universal gate set and have stable excitations akin to those of the two-dimensional Ising model within the color code family of models one needs $D = 6$ (Bombin *et al.*, 2013). Together with preparation and measurement in both the Pauli-X and Pauli-Z bases, the six-dimensional color code is compatible with the transversal application of the $\pi/8$ gate, and the controlled-not gate, which gives rise to universal fault-tolerant quantum computation.

Six dimensions are by no means a tight lower bound on the system dimensionality where both of these features coincide. Instead, this result is to be understood as a first estimate on the lowest spatial dimensionality that is to be improved upon. The result is obtained for the restricted case of models of commuting two-level physical systems. One may indeed be able to reduce the discovered critical dimension by considering many-body systems composed of higher-dimensional spins or fermionic degrees of freedom (Bombin *et al.*, 2013). Moreover, this result is restricted to models that give rise to fault-tolerant quantum computation by transversal gates. Indeed, quantum coding theory has shown that universal transversal operations are known to be incompatible with stabilizer error correction (Eastin and Knill, 2009; Zeng, Cross, and Chuang, 2011; Anderson and Jochym-O'Connor, 2014). The reader may

question this remark as to how the six-dimensional color code achieves a universal set of operations given known restrictions on universal transversal gate sets. Indeed, its transversal gate set is not truly universal, but, as pointed out earlier, its universal set of operations are completed by the ability to prepare and measure in both the Pauli-*Z* and Pauli-*X* bases. We finally remark that we may find stable low-dimensional systems with universal fault-tolerant operations by considering different types of fault-tolerant operations other than transversal gates.

VI. INTERACTING ANYON MODELS

As discussed in Sec. II.E, the syndrome of two-dimensional topological stabilizer codes can be interpreted in terms of pointlike anyonic quasiparticles. This leads to some favorable properties, such as the simple structure of the stabilizer space and the intuitive means by which the syndrome may be decoded. However, as explained in Sec. IV, it is precisely this pointlike nature that prevents the realization of self-correcting memories in two dimensions.

One path toward self-correction is to consider models with all the advantages of two-dimensional topological codes, but which are nevertheless able to realize self-correcting behavior. Interacting anyon models have been proposed to this end. All consider coupling a toric or planar code (Bravyi and Kitaev, 1998; Dennis *et al.*, 2002), the toric code Hamiltonian defined with open boundary conditions, to an external system, and then using this to mediate interactions between the anyons. These interactions change the energy landscape of the anyons and can lead to models with diverging energy barriers. Here we review the different types of anyonic interactions. We go on to review proposals to generate them.

A. Forms of anyonic interaction

To understand interacting anyon models, it is useful to first make a distinction between stabilizers and projective anyon number operators. The toric code stabilizers, introduced in Sec. II.D, are known as star and plaquette operators, denoted as A_v and B_p, respectively. Star and plaquette operators have eigenvalues ± 1. A state where a vertex v or a plaquette p of the toric code lattice holds an anyon lies in the -1 eigenspace of its corresponding stabilizer. Vacuum states, where there is no anyon on a given plaquette or vertex, lie within the $+1$ eigenspace of all the stabilizers.

By convention the stabilizer space corresponds to the $+1$ eigenspace of all stabilizers. We can use stabilizer operators to define projectors onto the common -1 eigenspace of the code and hence onto anyon states, such that

$$n_v = \tfrac{1}{2}(\mathbb{1} - A_v), \qquad n_p = \tfrac{1}{2}(\mathbb{1} - B_p). \qquad (63)$$

We call these projectors anyon number operators, as their eigenvalues are the anyonic occupation of vertex v or plaquette p.

For local Hamiltonians, the replacements $A_v \to -n_v$ and $B_p \to -n_p$ define an equivalent Hamiltonian, up to a constant shift in energy. However, interacting Hamiltonians that use n_j projectors lead to different physics. To illustrate this point, we

consider the example of two different Hamiltonians that describe interactions between a single pair of vertices, which we index 1 and 2. These Hamiltonians are

$$
\begin{aligned}
H_s &= -A_1 - A_2 - \tfrac{1}{2}A_1 A_2, \\
H_n &= -A_1 - A_2 + n_1 n_2.
\end{aligned}
\qquad (64)
$$

The first two terms ensure that the ground state is that of anyonic vacuum. The remaining term is an interaction.

The A_1 and A_2 terms contribute an energy penalty of 1 for each anyon present in both cases. The $A_1 A_2/2$ term contributes an energy penalty of 1 when there is a single anyon, but nothing when there are two. The $n_1 n_2$ term contributes nothing for a single anyon, but an energy penalty of 1 for a pair of anyons.

The different behaviors of H_s and H_n lead to different interpretations about what form the interactions take. For H_n the $n_1 n_2$ contributes only when multiple anyons are present. We therefore call it an anyon-anyon interaction. The $A_1 A_2/2$ term of Hamiltonian H_s contributes only when there is both an anyon on one vertex and vacuum on the other, so we call it an anyon-vacuum interaction. In what follows we consider both of these interaction types extended over the entire lattice.

B. Interacting anyon Hamiltonians

The first proposals for interacting anyon models considered anyon-anyon interactions of the form

$$H_{\mathrm{AA}} = \Delta H_{\mathrm{toric}} + V \sum_k \sum_{k' \neq k} n_k n_{k'} U(r_{kk'}), \qquad (65)$$

with H_{toric} defined in Eq. (15), values Δ and V are arbitrary coupling constants, indices k denote all stabilizers, both vertices and plaquettes, and $r_{kk'}$ is the Euclidean distance between k and k'. The first term here is the standard stabilizer Hamiltonian and the second is the anyon-anyon interaction.

1. Logarithmic potential

Hamma, Castelnovo, and Chamon (2009) proposed the following attractive potential:

$$U(r_{kk'}) = \ln r_{kk'}. \qquad (66)$$

Note that this diverges with distance, leading to a diverging energy barrier of $\mathcal{O}(\ln L)$. This potential has a confining effect for temperatures of $T < V/2$. In this regime, all anyons present in the code will typically be within an $\mathcal{O}(1)$ distance of each other. A logical error is caused when a single anyon breaks free and winds around the torus. This corresponds to a random walk in an $\mathcal{O}(\ln r_{kk'})$ potential formed by the other anyons, from which the typical coherence time $\tau = \mathcal{O}(L^{V\beta})$ can be found. This diverges polynomially with L and has an exponent that increases as temperature decreases.

For higher temperatures the model continues to have a diverging coherence time. This is due to a different mechanism, described in more detail later. The scaling in this case is $\tau = \mathcal{O}(L^2)$ (Wootton, 2013). Note that, although this remains polynomial, the exponent no longer depends on temperature.

This analysis applies only to the toric code. For the planar code (Bravyi and Kitaev, 1998; Dennis *et al.*, 2002) a single anyon can be created at a boundary. Since this has no other anyons to confine it, it does not experience a $\tau = \mathcal{O}(L^{V\beta})$ coherence time for any temperature. Instead the coherence time scales as $\tau = \mathcal{O}(L^2)$ for all temperatures.

An interaction of this form can be mediated by coupling the anyons to a two-dimensional lattice of hopping bosons. Only local couplings are needed to produce the interaction (66). However, the coupling strength must diverge with system size to realize the potential. This violates a requirement for a self-correcting memory, that of bounded interactions. Nevertheless, note that these diverging couplings are not directly responsible for the increasing lifetime, since the physical energy scales remain finite. Moreover, the model needs to be fine-tuned. Small perturbations in the hopping couplings of the bosons can cause the long-range interactions to become short range and thus stop the effectiveness of the model.

2. Power-law potential

In Chesi, Röthlisberger, and Loss (2010) the interactions have the following repulsive potential:

$$U(r_{kk'}) = r_{kk'}^{-\alpha}, \qquad \alpha \geq 0. \tag{67}$$

At first glance this would appear to be ineffective. The potential does not diverge with distance and in fact it decays in general such that the energy barrier for the creation and separation of a pair is finite. Furthermore, the potential is repulsive, making anyons less likely to annihilate once created than in the noninteracting case. Nevertheless, when the interactions are sufficiently long range $\alpha < 2$ they have a strong beneficial effect.

The simplest case to consider is that of $\alpha = 0$, where the potential does not decay over distance but remains constant. The energy of the system does not depend on the positions of anyons in this case, only their number N. Every anyon repels every other one with an energy V. The energy above the ground state is then

$$\varepsilon_N = N\Delta + \frac{V}{2}N(N-1). \tag{68}$$

Note that this grows quadratically with N, rather than simply linearly as in the noninteracting case. A similar superlinear scaling of energy with anyon number occurs for other $\alpha < 2$ as well as the logarithmic potential.

By using simple arguments, it is possible to find a lower bound on the coherence time for the $\alpha = 0$ case. Since we can realistically expect only a finite energy density, i.e., states with $\varepsilon_N = \mathcal{O}(L^2)$, the system is limited to states with $N = \mathcal{O}(L)$ anyons at most. This corresponds to an infinitely sparse anyon configuration, with a distance of $\mathcal{O}(\sqrt{L})$ between each pair of anyons created by the thermal noise. The time required to cause a logical error is then no less than that required for random walks over this length scale, and so $\tau = \mathcal{O}(L)$.

Using a more careful treatment it can be shown that the number of anyons is suppressed further than the argument

implies. This leads to a longer coherence time which, for general $\alpha < 2$, scales like $\tau = \mathcal{O}(L^{2-\alpha})$. This is a polynomial scaling that is quadratic in the best case. This effect is also responsible for the $\tau = \mathcal{O}(L^2)$ scaling of coherence time for the logarithmic potential we met earlier.

The power-law potential can be mediated by interaction with cavity modes. This is a nonlocal interaction, but is nevertheless reasonable up to a cutoff system size. The self-correcting behavior is therefore not truly scalable, as we require. However, it may still be possible to achieve system sizes that are useful in practice.

C. Anyon-vacuum interactions

Recently interacting anyon model proposals have focused on engineering couplings between anyons and the vacuum rather than interactions between anyons. It is found that this approach allows significantly more powerful self-correction. The first studies of such models are found in Pedrocchi, Chesi, and Loss (2011) and Hutter *et al.* (2012).

These proposals consider Hamiltonians of the form

$$H_{\text{AV}} = \Delta H_{\text{toric}} - V\sum_k \sum_{k' \neq k} S_k S_{k'} r_{kk'}^{-\alpha}. \tag{69}$$

Here S_k are the stabilizer operators, either A_v or B_p depending on whether each k is a vertex or a plaquette. The only difference between this Hamiltonian and that of the power-law potential is the substitution of anyon number projectors n_k with stabilizer operators S_k.

To compare H_{AV} with that for the anyon-anyon case H_{AA}, we can rewrite it in terms of anyon number operators. This yields

$$H_{\text{AV}} = \mu(L)\sum_k n_k - 4V\sum_k \sum_{k' \neq k} n_k n_{k'} r_{kk'}^{-\alpha} + \text{const.} \tag{70}$$

The first term here is the effective anyon gap, the energy that an anyon must overcome in order to be created

$$\mu_k(L) = \Delta + 4\sum_{k'}(1 - \delta_{k,k'})r_{kk'}^{-\alpha}. \tag{71}$$

The large energy penalty for creation is due to the anyons being repelled by the majority of plaquettes and vertices, which are in the vacuum state. The second term in Hamiltonian (70) is an attractive and nondivergent anyon-anyon attraction.

The effects of thermal errors are suppressed first by the anyon gap, which significantly slows the creation of anyons. Once anyons have been created, errors are further suppressed by the attractive potential that favors reannihilation. Let us focus on the effects of the anyon gap.

To suffer a logical error, an error must first occur on a single qubit. This will create at least n anyons, where $n = 1$ for the planar code, for the case where a qubit is created on the boundary, and $n = 2$ for the toric code. These n anyons will feel a repulsion from all the $\mathcal{O}(L^2)$ plaquettes and vertices on which there is vacuum. This results in the energy gap

$\mu(L) = \mathcal{O}(L^{2-\alpha})$ for each anyon when $\alpha < 2$. The typical time before the first error occurs is then $e^{\beta n \mu(L)}$.

It is this time scale that dominates the lifetime of the code. Other factors must be taken into account to fully deduce the lifetime: the number of qubits on which it is possible for an error to occur; the probability that a pair, once created, will cause a logical error; and the time required for an anyon to diffuse across the code. However, these will contribute factors that are polynomial in L at most. They are insignificant in comparison to the exponential time scale to overcome the anyon gap. As such, we can simply say that $\tau = e^{\beta n \mu(L)} = e^{\beta \mathcal{O}(L^{2-\alpha})}$. Note that the higher n for the toric code will result in asymptotically longer lifetimes than the planar code, assuming all else is equal.

The main difference between proposals with this interaction is the physical system used to mediate it and the value of α that is achieved. Pedrocchi, Chesi, and Loss (2011) and Hutter *et al.* (2012) achieved the optimal case of $\alpha = 0$ by coupling to cavity modes. However, as noted, this prevents the model from achieving the scalability that we require. A Hamiltonian simulation of this case was considered by Becker *et al.* (2013).

Pedrocchi *et al.* (2013) mediated the interaction by coupling to a three-dimensional lattice of hopping bosons, as shown in Fig. 12. The resulting interacting model corresponds to that of $\alpha = 1$. This gives rise to an energy barrier of $\mu(L) = \mathcal{O}(L)$ and coherence time $\tau = e^{\beta \mathcal{O}(L)}$. This model uses only local, bounded strength and constant weight interactions and yet is able to preserve the quantum information for an exponentially long time.

A related model was also proposed by Hutter *et al.* (2014). Here the role of the bosons is played by magnons in a three-dimensional ferromagnet. The code is simply coupled to the spins of the ferromagnet, avoiding the need for unbounded operators. Perturbative gadgets (Bravyi *et al.*, 2008; Bravyi, DiVincenzo, and Loss, 2011) are used in order to realize the entire Hamiltonian using only local two-body interactions on a three-dimensional system of spin-half particles.

The effective interacting anyon model in the case of the ferromagnet is the same as for the three-dimensional boson
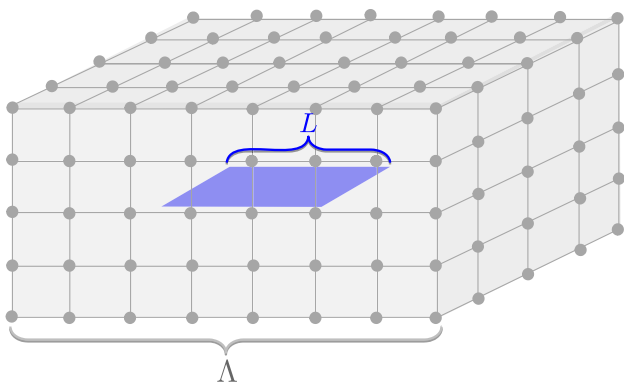

FIG. 12. Schematic of the planar code coupled to a bosonic bath. Anyon-vacuum interactions in an $L \times L$ code can be induced by embedding them in a Λ^3 lattice of hopping bosons for $\Lambda > L$. The stabilizers of the code locally couple to the bosonic lattice. The interactions are mediated by the low-energy collective excitations of the bosons.

lattice. However, it is expected to be valid only in the regime for which the coupling of the code to a thermal bath is much weaker than the coupling to the ferromagnet. The Monte Carlo simulation of thermal noise described in Sec. III.E therefore cannot be used. This approximates the true dynamics by allowing the thermal bath to instantaneously transfer large energies to the system and so is unable to capture the subtle effects arising from the full time evolution of the system-bath interaction. It is in these effects that we expect to observe self-correcting behavior. Since a study of this time evolution seems intractable, the full extent of the self-correcting behavior in this model is not known. This model can also be adapted to protect topological systems composed of superconducting qubits for which errors correspond to infinitely weak coupling to an infinite-temperature bath (Kapit, Chalker, and Simon, 2014).

D. Open questions in interacting anyon models

As discussed, interacting anyon models are capable of impressive memory times which grow polynomially or even exponentially with system size. However, due to the difficulties in solving the proposed interacting models, further study is required to better understand potential challenges we may need to overcome to realize effective long-range interactions. In this section we outline some of the important directions of study that we must follow to understand the feasibility of realizing a self-correcting quantum memory by means of interacting anyons.

One such direction concerns the effects of local perturbations on interacting anyon models. The nonlocal terms of an interacting anyon Hamiltonian can cause local pertubations to affect the system nonlocally. Any perturbation that gaps the bosons, for example, will cause the anyonic interactions to become short range. This will induce a finite cutoff length scale, beyond which the lifetime no longer increases with system size. This was discussed for the ferromagnet based model in Hutter *et al.* (2014) and elaborated upon for the bosonic model in Landon-Cardinal *et al.* (2015).

Such problems with perturbations are not unique to proposals for quantum memories. The same is true for the ferromagnetic systems currently used as classical memories. Perturbations in these systems limit the magnetic susceptibility, inducing an equivalent finite cutoff. However, this has not constrained the development of classical computation. It is therefore important to determine how large the cutoff will be for quantum memories when realistic perturbations are considered. Whether or not it is large enough to allow for large-scale quantum computation is a crucial test of these proposals and remains an open area of study. Additionally, possible avenues toward the discovery of interacting anyon models that are stable against perturbations have also been discussed (Landon-Cardinal *et al.*, 2015). Concrete examples, however, remain to be found.

Another caveat of the interacting anyon proposals concerns the thermal bath used for their analysis. In all of the cases discussed in this review, it was assumed that the thermal bath will act locally on the anyonic system. However, it may be that thermalization will occur in a more complex manner, with the mediating system allowing thermal errors to become long

range. Analysis of the self-correcting properties with realistic thermalization dynamics therefore remains an open area of study.

E. Finite-temperature topological order

It is important to determine if a system can maintain topological order at nonzero temperature. We might expect such a system to be stable against local perturbations which is a property that we require of a self-correcting quantum memory. For many systems, topological order vanishes at nonzero temperature (Nussinov and Ortiz, 2008, 2009a, 2009b; Chesi *et al.*, 2010). Indeed, it was shown by Hastings that local commuting Hamiltonians in two dimensions are incompatible with a condition for topological order at finite temperature (Hastings, 2011). Some models, however, demonstrate a finite-temperature phase transition between a phase with topological order and a disordered phase.

The relationship between topological order and self-correction is not well understood in the general case (Yoshida, 2011). So far, we have considered only specific models to learn the correspondence between topological order and self-correction; see, for instance, Mazáč and Hamma (2012) where topological entanglement entropy and the coherence times of the toric code are studied on lattices of varying dimensionality. For the cases studied there is a strong coincidence between these two features. The interacting anyon models are also shown in Wootton (2013) to provide an interesting perspective on this problem which we discuss in this section.

The signature of topological correlations in systems such as the toric code are loop correlations. For finite-temperature systems these can be found using the topological entropy (Hamma, Ionicioiu, and Zanardi, 2005a; Kitaev and Preskill, 2006; Levin and Wen, 2006), topological mutual information (Iblisdir *et al.*, 2009), or anyonic topological entropy (Wootton, 2013). In all cases one must consider a region of the system whose size is on the order of the system size which can be arbitrarily large. If topological correlations are detected for large regions, the state is said to be topologically ordered.

For power-law anyon-anyon interactions, any finite-temperature thermal state contains a diverging number of anyons (Chesi, Röthlisberger, and Loss, 2010). These anyons will also be deconfined due to the repulsive nature of the interactions. Such a diverging number of delocalized anyons ensures that topological order according to the above definition is not present for any finite temperature. However, the interactions are still known to support self-correction with a polynomial lifetime. This fact, along with the $\mathcal{O}(1)$ energy barrier, makes these models an interesting exception to widely held opinions about what is required for self-correction.

For the case of the logarithmic anyon-anyon interaction, it is found that the topological order persists up to a finite temperature of $T_c = V/2$. It therefore corresponds exactly to the confined anyon phase for which the lifetime is $\tau = \mathcal{O}(L^{V\beta})$. Beyond this temperature the topological order is no longer present, although the system is still self-correcting with a lifetime of $\tau = \mathcal{O}(L^2)$. The phase transition does not have the effect of destroying the self-correction, as we might expect, but it does alter the scaling of the lifetime.

To regain some semblance of our intuition that topological order is required at finite temperature for self-correction, we can redefine what we mean by topological order. Instead of simply considering whether topological correlations can be detected for regions of an arbitrarily large size, we can determine the exact range of these correlations. This means considering regions of different sizes, calculating how the topological correlations decay as the size is increased, and then determining a correlation length to quantify this. This correlation length is denoted λ.

Even when a system is not topologically ordered, it is possible for topological correlations to be present on an $\mathcal{O}(1)$ length scale. The case of $\lambda = \mathcal{O}(1)$ therefore corresponds to topologically trivial states. Standard topological order requires that the correlations do not decay at all, even up to the linear system size L. This corresponds to a super-extensive $\lambda > \mathcal{O}(L)$. In between these extremes there exists the possibility for the range λ to increase with system size, but just not as quickly as is required for standard topological order. We refer to such states as weakly topologically ordered.

Studying the interacting anyon models from this perspective, it has been shown that all models at all temperatures are in either a standard or weakly topologically ordered phase (Wootton, 2013). All transitions between these two types of topological order correspond to a change in the way that the lifetime scales with system size. All known self-correcting memories correspond to phases that are either topologically ordered or weakly topologically ordered. As such, some relationship between finite-temperature topological order and self-correction does persist. However, it would be interesting to determine whether counter examples exist even for the case of this weaker relationship.

VII. COMMUTING THREE-DIMENSIONAL MODELS

With limitations challenging the construction of finite-temperature quantum memories with commuting two-dimensional Hamiltonians, it is exciting to consider three-dimensional models. Indeed, as discussed in Sec. IV there are still no-go results toward the feasibility of a finite-temperature quantum memory in three dimensions. However, these results are not as restrictive as their two-dimensional counterparts. Recently proposed models have shown positive progress, which, together with supporting numerical data offer promise for the discovery of good quantum memories at finite temperature. In this section we provide an overview of the positive results found in three-dimensional models. We first review the concept of partial self-correction: a new paradigm for macroscopic coherence time scaling that has emerged from the study of three-dimensional models. In Sec. VII.B we study the cubic code model, a quantum system demonstrating partial self-correction. In Secs. VII.C and VII.D we review other three-dimensional proposals that break translational invariance to achieve phenomena potentially important for self-correction in quantum systems.

A. Partial self-correction

A new phenomenon to develop from the study of three-dimensional systems is that known as *partial self-correction*.

Partially self-correcting models are notable for polynomial coherence-time scaling with system size up to some cutoff size that depends on temperature. Moreover, they exhibit super-exponential inverse-temperature scaling. These features of known partially self-correcting models are attributed to its energy barrier, which grows logarithmically with the size of an error incident to a memory.

Partially self-correcting models were discovered independently by both Haah (2011) and Castelnovo and Chamon (2012) following remarkably different methods. Haah (2011) exhaustively searched over all translationally invariant stabilizer models on a cubic lattice with one or two qubits on each vertex of the lattice to find models that satisfy the "no-strings" condition, as defined rigorously by Haah (2011). Broadly speaking, the no-strings condition is a property of the excitation structure of a commuting Pauli-Hamiltonian, whereby a cluster of nontrivial excitations cannot be transported across a lattice over an arbitrary distance without introducing new excitations to the lattice. Models that satisfy the no-strings condition are thus expected to provide good protection against thermal interactions due to the large energy cost that is required to introduce logical errors to the lattice via local operations. Indeed, it was later proved that models satisfying the no-strings condition must have an energy barrier that scales at least logarithmically with the size of the system (Bravyi and Haah, 2011b). Independent of the work of Haah, Castelnovo and Chamon (2012) looked to find a quantum generalization of known low-dimensional classical models with nontrivial energy barriers between ground states which are well known in the context of glassy systems (Newman and Moore, 1999; Garrahan and Newman, 2000). The derived generalization is described locally in three dimensions. Similar models are studied in further generality by Kim (2012).

We now give a heuristic analysis explaining partial self-correction where we assume Arrhenius's law, Eq. (29). Known partial self-correcting memories are characterized by an energy barrier that grows logarithmically with the size of an error ξ. The energy ε that is required for an error to increase in size to occupy a volume of lattice of diameter ξ must thus be at least

$$\varepsilon \sim \kappa \Delta \log \xi, \tag{72}$$

where κ is a positive constant and Δ is the gap of the model. Given the typical finite-temperature noise analysis we described in Sec. III, we assume that errors are created with an average separation that scales with β like $\Lambda \sim e^{\beta \Delta / D}$, where D is the dimensionality of the system. Given then that we require $\xi \sim \Lambda$ for the memory to decohere, we arrive at the typical excitation energy of the model as a function of β. Namely, we obtain $\varepsilon \sim \kappa \Delta^2 \beta / D$ at the point of decoherence. Applying this to Arrhenius's law, Eq. (29), we obtain

$$\tau \sim e^{\kappa \Delta^2 \beta^2 / D}. \tag{73}$$

We can follow a similar analysis to study small system sizes such that $L \lesssim \Lambda$. In this case a diffusing error must attain energy $\varepsilon \sim \kappa \Delta \log L$. Once again, applying this expression to

Arrhenius's law it follows that partially self-correcting quantum memories have a coherence time that grows polynomially in system size

$$\tau \sim L^{\kappa \Delta \beta}, \tag{74}$$

whose exponent is linear in β. This scaling is effective up to some cutoff $L_{\text{opt}} \sim \Lambda$. We thus obtain a cutoff $L_{\text{opt}} \sim e^{c\beta}$ that grows exponentially in β for positive constant c (Bravyi and Haah, 2013).

An interesting feature of the known partially self-correcting models is that they do not have a constant ground-state degeneracy, which is outside the assumptions of the no-go theorem of Yoshida (2011). As an aside, it is interesting to consider that the known partially self-correcting models are not purely topological models. While the ground space of these systems is topologically ordered in the sense that the degenerate ground states of the model cannot be locally distinguished, their ground-state degeneracy still depends on the microscopic physics of the model. A similar model whose ground-state degeneracy depends on microscopic details (Bravyi, Leemhuis, and Terhal, 2011) is introduced by Chamon (2005). We remark, however, that unlike the cubic code, the model due to Chamon has a constant energy barrier and is well understood not to give rise to self-correcting properties (Chamon, 2005; Castelnovo and Chamon, 2012; Nussinov, Ortiz, and Cobanera, 2012; Temme, 2014). We finally remark that the study of exotic partially self-correcting systems has led to new classifications of systems under the context of fractal topological quantum field theories (Yoshida, 2013; Haah, 2014).

In the following section we review and reproduce previously obtained numerics of the rigorously studied cubic code model. We remark that the fragile glassy model introduced by Castelnovo and Chamon (2012) is expected to behave in a phenomenologically equivalent way to the cubic code model.

B. The cubic code

The cubic code model (Haah, 2011) is defined on a three-dimensional lattice of $L \times L \times L$ vertices, where two qubits lie on each vertex of the lattice. Associated with each of the unit cubes of the lattice, indexed j, we have two stabilizers S_j^X and S_j^Z shown in red (left) and blue (right) in Fig. 13, respectively. We then write the Hamiltonian

$$H_{\text{cubic}} = -\frac{1}{2} \sum_j (S_j^X + S_j^Z). \tag{75}$$

We take constant interaction strength $1/2$ such that excitations have unit energy cost. The model has a nontrivial ground-state degeneracy that varies with L. The ground-state degeneracy is studied in detail by Haah (2013) using the language of commuting free modules. For simplicity, we consider only lattices of size $L < 200$ that do not have factors 2, 15, or 63. All of these system sizes have a fourfold ground-state degeneracy (Bravyi and Haah, 2011a). We do not discuss the complex fractal structure of the logical operators of the

model here, but we refer the interested reader to Haah (2013). We point out, however, that for lattices of odd L, we find two logical operators by taking $\bar{X}_1 = \prod_{k \in \mathcal{L}} X_k$, $\bar{Z}_1 = \prod_{k \in \mathcal{L}} Z_k$, $\bar{X}_2 = \prod_{k \in \mathcal{R}} X_k$, and $\bar{Z}_2 = \prod_{k \in \mathcal{R}} Z_k$, where $\mathcal{L}$ and $\mathcal{R}$ denote the subset of all the left and right qubits of each vertex, respectively. It is easily checked that these operators satisfy a suitable algebra for the logical qubits of the model. These of course are not the minimum-weight logical operators (Bravyi and Terhal, 2009). However, we find these logical operators particularly convenient for numerical simulations.

1. Excitations of the cubic code

The excitations of the cubic code have a more complicated structure to those of the two-dimensional models considered in earlier sections. Indeed, the model was designed such that its excitations are created by operators that satisfy the no-strings rule (Haah, 2011) and instead have a fractal-like structure. In Fig. 13 we can observe a symmetry over the S_j^X and S_j^Z stabilizers, such that both Pauli-X and Pauli-Z–type errors act with equal effect on the S_j^Z and S_j^X stabilizers of the model. Analysis of only Pauli-X–type errors therefore serves for a sufficient study of the excitations of the cubic code.

Pauli operators acting on single qubits of the lattice in the ground state create four localized excitations to the dual lattice of the model, shown in Figs. 14(a) and 14(b). We mark the stabilizers violated by the error by shaded (yellow) points on the vertices of the dual lattice, where vertices of the dual lattice correspond to fundamental cubes of the primal lattice. Similar to the excitations of the stringlike models we already discussed, these particles are delocalized. Also, the excitations are their own antiparticles, which are transported by applying additional error operators that annihilate excitations and create additional excitations at other locations on the lattice. In this way, it is possible to delocalize these excitations over arbitrary distances. We show an error configuration in Fig. 14(c) where four fundamental excitations have delocalized over two lattice spacings.

An important distinction between the excitations of the cubic code model and excitations in two-dimensional models is that the delocalization of these particles cannot be achieved using stringlike operators. Instead, if we want to delocalize the excitations of the cubic code model over arbitrary distances, we have to use transport operators that have a fractal-like support. As a finite-temperature noise model will only apply transport operators to the lattice via local single-qubit operations, such operators are only achieved by temporarily increasing the energy of the error configuration, which
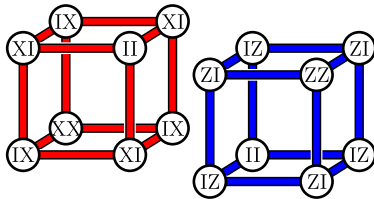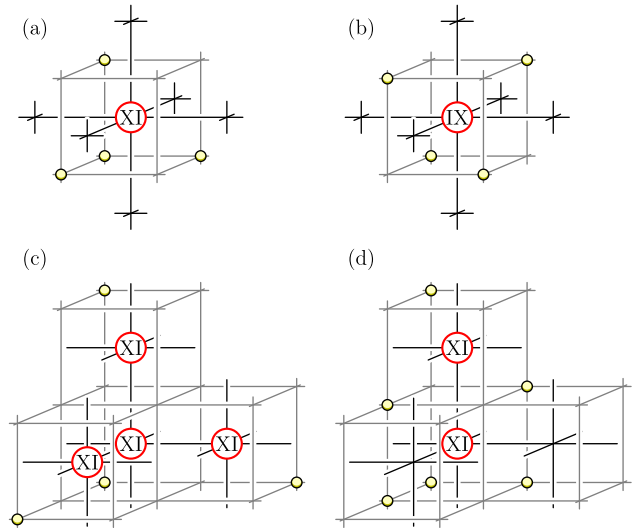
FIG. 14. The excitations of the cubic code. The lattice supporting the cubic code is shown in black where two qubits lie on the vertices of the lattice. The dual lattice that supports excitations is shown in gray. (a), (b) The excitations generated by an *XI* error and an *IX* error acting on a single two-qubit vertex, respectively. Both errors generate four pointlike excitations, marked by shaded (yellow) points on the vertices of the dual lattice. (c) The error configuration that creates four excitations delocalized over two lattice spacings. (d) An example of a high-energy intermediate error configuration that must be achieved to delocalize four excitations over a long distance as in (c).

reduces the propagation of excitations. In Fig. 14(d), we show an intermediate error configuration necessary to delocalize excitations over two lattice sites.

The error configuration creates six excitations, which increases the energy of the system. Models where excitations are not propagated by stringlike operators satisfy the no-strings rule. This is an important concept for partial self-correction. An extensive program of analytical study from Bravyi and Haah (2011b) has proved that models satisfying the no-strings rule necessarily have at least a logarithmic energy barrier. Further work in this program of research numerically showed that the cubic code model behaves as a partially self-correcting memory with a logarithmic energy barrier (Bravyi and Haah, 2013). Beyond the study of self-correcting memories, the cubic code is also noteworthy from the point of view of localization. Kim and Haah (2016) showed that the glassy nature of the cubic code localizes the excitations of the model. This is particularly interesting since localization is typically attributed to disorder (Anderson, 1958). Localization in the cubic code however is achieved with a frustration-free Hamiltonian with uniform interactions.

In the remainder of this section we numerically simulate the cubic code at finite temperature to demonstrate its partial self-correcting behavior.

FIG. 13. The stabilizers of the cubic code. The model is defined on a cubic lattice with two qubits on each vertex. Stabilizers denoted S_j^X and S_j^Z are shown on the left and right, respectively.

2. Numerical simulations

In this section we simulate the cubic code coupled to a finite-temperature environment using the numerical Monte Carlo methods described in Sec. III.E. We reproduce

the phenomenological behavior demonstrated by Bravyi and Haah (2013). We remark that the results we present differ from those given by Bravyi and Haah (2013) due to the choice of rate equation used in the simulation. Where we use the rate equation discussed in Sec. III.A, Bravyi and Haah use the rate equation of Bortz, Kalos, and Lebowitz (1975). Both rate equations satisfy the detailed balance given in Eq. (19) and ultimately will reproduce the same physics up to some variation in the obtained constant factors.

We simulate Hamiltonian (75) under rate equation (18). To obtain coherence times, we periodically attempt to decode the state of the evolving lattice using a variant of the clustering decoder described in Appendix A. The first time at which the decoder fails gives the coherence time of the sample. We find the coherence time by averaging over N samples. Errors are determined by taking the standard deviation of the samples, divided by $\sqrt{N}$.

Identifying coherence times for the cubic code using numerical simulations is particularly challenging due to the glassy nature of the model (Chamon, 2005; Castelnovo and Chamon, 2012). Specifically, as the system evolves toward the equilibrium state, the simulation frequently finds states that are local minima of the considered Hamiltonian. At the computational level we must simulate many events to escape these metastable configurations, which we find to be numerically intensive. To overcome this, for the small system sizes we study, we find that decoding with very high frequency reduces the number of error events that we must simulate as the state decoheres. We attempt to decode the system at time intervals $\sim 10^{-10} e^{4\beta}$. The time units $e^{4\beta}$ are natural as this is the frequency at which four excitations are created from the vacuum which then mutually annihilate themselves shortly afterward with high probability. This behavior is typical in the small system size and low-temperature regime.

To identify partial self-correction in the cubic code we plot the coherence times as a function of L in Fig. 15. Here we consider many different temperatures to identify polynomial coherence time scaling with system size whose exponent depends on β. We plot the gradients of the linear fittings, shown in the inset of Fig. 15. The gradients we obtain show good agreement with polynomial coherence time scaling whose exponent grows linearly with β, as we expect for partial self-correction, derived in Eq. (74). We also plot the optimal coherence time as a function of system size for a given β, as shown in Fig. 16. We identify superexponential inverse-temperature scaling, as we expect for a partially self-correcting model due to Eq. (73).

We use the fittings we obtain from the presented numerical data to obtain Δ and κ of Eqs. (73) and (74) for the cubic code model

$$\Delta_{\text{CC}} = 2.0, \qquad \kappa_{\text{CC}} = 0.79, \qquad (76)$$

thus identifying the partial self-correcting behavior described in the previous section. It is interesting that diverging coherence times at low temperatures are achieved here via a glassy mechanism (Castelnovo and Chamon, 2012), and not with some ordered phase of matter such as those we considered in Secs. III.D and V. The glassy nature of the
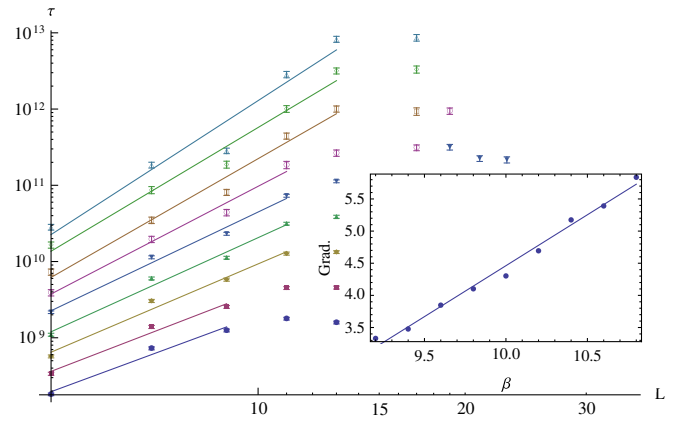

FIG. 15. Numerically evaluated coherence times τ for the cubic code shown as a function of system size L. We show separate lines for $\beta = 9.2, 9.4, \ldots, 10.8$, where $\beta = 9.2$ is shown by the lower dark blue line and $\beta = 10.8$ is shown by the light blue line at the top. The inset shows the gradients of each of the lines in the main figure which, in this temperature regime, grow like $\tau \sim L^{1.58\beta - 11.38}$.

model may introduce new difficulties in encoding information to the cubic code since cooling the system to its ground space will be very slow (Chamon, 2005). Instead, we might consider some manual method of state preparation by measurement or otherwise. Work in this direction has been conducted by Lodyga *et al.* (2015).

Haah (2013) showed that by imposing translational invariance on three-dimensional commuting Pauli-Hamiltonians we cannot expect to find a system where the energy barrier scales better than logarithmically with system size. In the remainder of this section we consider commuting Pauli-Hamiltonian models that surpass the result of Haah (2013) by breaking translational invariance. The models of interest are the welded three-dimensional toric code (Michnicki, 2014) and embeddable-fractal product codes (Brell, 2016).

C. The welded toric code

The welded toric code, due to Michnicki (2012, 2014), is the first explicit example of a three-dimensional commuting Pauli-Hamiltonian with a power-law energy barrier. Remarkably, the model surpasses the no-go result of Haah (2013) by breaking the translational invariance assumption that is required to complete the theorem.

The model is found using an idea called *welding*, described by Michnicki (2012). Welding gives a procedure to combine stabilizer codes. The advantage of welding codes is that logical operators are also combined nontrivially over a weld. We follow the exposition of Michnicki showing how he arrived at the welded toric code.

The welded toric code is achieved by welding a macroscopic number of copies of the three-dimensional toric code (Hamma, Zanardi, and Wen, 2005). The three-dimensional toric code is defined on a cubic lattice with qubits on the lattice edges. The model has two types of stabilizers, vertex operators, and face operators, as shown in Fig. 17(a). Vertex operators are six-body operators of Pauli-X operators supported on the edges incident to a vertex. Face operators are
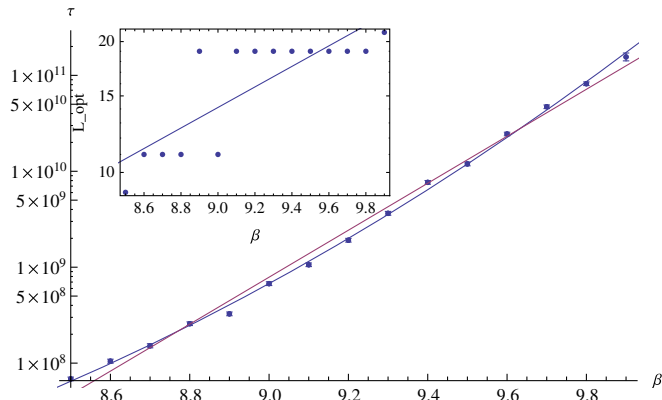
FIG. 16. Coherence-time data τ for the cubic code plotted as a function of inverse temperature β. The plotted data are found using system size $L = L_{\text{opt}}$ for each value of β. We also show L_{opt} as a function of β in the inset. We fit the data shown in the main figure to the function $\tau \sim e^{1.05\beta^2 - 13.7\beta + 58.5}$. We compare this fitting to a fitting based on Arhenius's law where we find $\tau \sim e^{5.6\beta - 30.3}$, also shown in the figure. The inset shows $L_{\text{opt}} \sim e^{0.54\beta - 2.2}$.

four-body Pauli-Z operators supported on edges that bound faces of the cubic lattice.

With a suitable choice of boundary conditions the three-dimensional toric code will support one logical qubit. The model has two types of boundaries, rough boundaries and smooth boundaries. We show a macroscopic picture of the code in Fig. 17(b), where we have rough boundaries on the top and bottom faces of the lattice. Later, we will refer to one copy of the three-dimensional toric code that encodes a single logical qubit as a *block*.

There are two types of logical operators, membrane logical operators and string logical operators, shown in red and blue, respectively, in Fig. 17(b). The membrane logical operators have dynamics akin to those of the two-dimensional Ising model and as such are stable below a critical temperature (Castelnovo and Chamon, 2007; Alicki *et al.*, 2010). However, stringlike logical errors introduce pointlike excitations whose creation and transport need only overcome a constant energy barrier. It is for this reason that a thermal noise model is able to introduce stringlike logical errors in constant time. The three-dimensional toric code therefore does not behave as a self-correcting quantum memory.

Michnicki surpasses the problem of stringlike logical operators using welding. He shows that it is possible to weld blocks of three-dimensional toric code along rough faces to generate a large energy barrier. We show an example of a weld in Fig. 17(c), where four copies of the three-dimensional toric code are welded along a common rough boundary. The diagram shows that the stringlike logical operator divides at the weld. A single pointlike excitation therefore cannot pass through a weld from one code block to another, but must instead split into $v - 1$ excitations where v is the valency of the weld, i.e., the number of code blocks that meet at the weld. Since these excitations must overcome an energy barrier to propagate across the weld, the diffusion of errors is suppressed.

Michnicki combines many three-dimensional toric code blocks in a latticelike structure with dimensions greater than 1 to generate a macroscopic energy barrier over the stringlike logical operators of the model. We show such a lattice of two dimensions in Fig. 17(d). Now, the stringlike logical operators of the composite parts of the code are combined into a coarse lattice whose vertices are welded faces and whose edges are code blocks. Interestingly, the model has interactions resembling high-dimensional Ising-like interactions across the welded boundaries, separated by three-dimensional blocks of toric code lattice. With this lattice configuration, thermal fluctuations must overcome a polynomial energy barrier in the number of welds of the macroscopic welded lattice to introduce a logical error to the model.
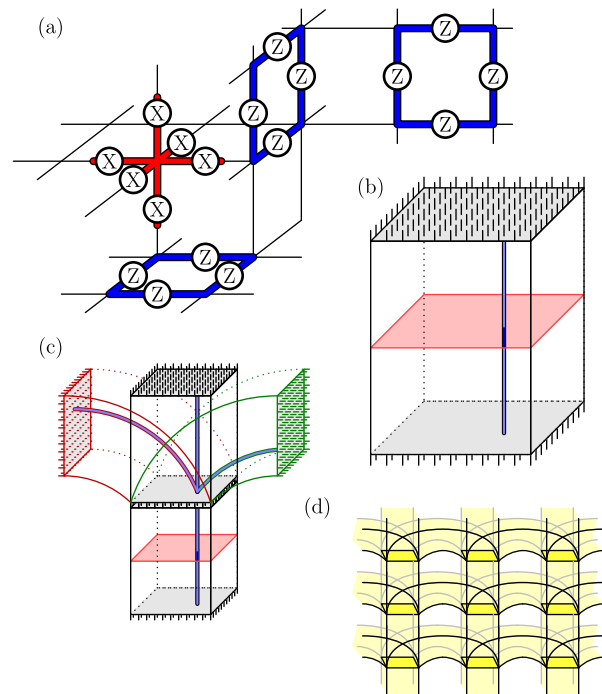

FIG. 17. The welded toric code. (a) The stabilizers of the three-dimensional toric code where qubits lie on the edges of the lattice. A vertex operator is shown in red in the top-left corner of the figure. Three examples of face operators are also shown on the lattice, each with a different orientation relative to the cubic lattice. Vertex operators are the tensor product of Pauli-X operators supported on all the edges adjacent to a vertex. Face operators are the tensor product of Pauli-Z operators supported on the edges that bound a face of the lattice. (b) The three-dimensional toric code with two disjoint rough faces. The model supports one encoded qubit described by a membrane logical operator, shown as a red horizontal plane, and a stringlike operator, shown by a vertical blue line that runs between the two disjoint rough edges. (c) A single weld between the rough faces of four blocks of three-dimensional toric code. We outline one block in red and one block in green where the blocks are overlapping. The stringlike logical operators of the original code blocks must now overcome the high-energy barrier at the interface between the welds. (d) Many blocks of three-dimensional toric code are welded into a lattice. In this welded configuration, the stringlike logical operators combine to give a two-dimensional logical operator that spans the lattice.

The membrane logical operator of a code block is not extensive with the number of welded interfaces of the lattice. In fact, the membrane logical operator can be supported on a single block of the code. To scale the power-law energy barrier correctly, the volume of the blocks must grow with the size of the welded lattice. Michnicki (2014) suggested scaling the volume of the block size polynomially with the number of welds in the lattice. It is with this point that we see the model breaks translational invariance; we can vary the size of the model over two different length scales, the block size, and the volume of the lattice of welds.

1. Excitations of the welded toric code

The welded toric code model shares many similarities to the Ising model, with the code blocks in the former serving a similar function as the nearest-neighbor ferromagnetic couplings, or "bonds," in the latter. We now discuss this analogy to find a better understanding of the finite-temperature behavior of the welded toric code.

We consider the simple case in which we do not know the full details of the excitation configuration within each code block. Instead we know only the total parity of the number of pointlike excitations that each block contains. Each code block then has two possible states: even or odd parity. These correspond to the two possible states of a bond of the Ising model: aligned and antialigned. In general we are able to determine the locations of all the excitations in a code block. However, since the excitations of the code blocks become disordered after a constant time, information regarding their locations becomes irrelevant. For this reason the total parity of each code block will be the only useful syndrome information after thermalization occurs. The error-correction procedure of the welded toric code is then equivalent to that of the corresponding Ising model.

The parity of a code block is changed only if a pointlike excitation moves through one of its welds, which occurs only if an error occurs on one of the qubits involved in a weld. We refer to a weld that has suffered an odd number of errors as "broken." Breaking a weld is equivalent to flipping the spin on a vertex of its corresponding Ising model. A region of broken welds (flipped spins) will be surrounded by a surface of odd parity codes (antialigned bonds). To perform a logical error, these surfaces must be removed by breaking all welds or flipping all spins, respectively.

Given this correspondence between the welded toric code and the Ising model, we can expect the former to inherit the exponential lifetime and finite-temperature phase transition of the latter. However, this is not the case, since the welds of the welded toric code experience temperature differently from the spins of its corresponding Ising model. We observe that due to the $\mathcal{O}(L^2)$ area of each weld, for code blocks of size $L \times L \times L$, the probability that any given weld is broken after the constant thermalization time converges exponentially to $1/2$ as $L \to \infty$. It follows from this that the probability of finding any code block in odd parity also quickly converges to $1/2$ as the model approaches the thermodynamic limit. For the Ising model at effective inverse temperature $\tilde{\beta}$, the probability of finding a bond in antialignment approaches $1/2$ only as $\tilde{\beta} \to 0$. Equating these probabilities it is clear that $\tilde{\beta}$ vanishes

as $L \to \infty$. Hence, the welded toric code at finite temperature corresponds to an Ising model with an effective temperature that diverges with system size, and so it does not fare well as a memory in the thermodynamic limit.

We next consider the low-temperature behavior of the welded code. To decay encoded information, the thermal environment must overcome a macroscopic number of welds. The most energetically challenging process for a weld to break is for an error to occur on a qubit involved in a weld. This event happens at a constant rate $\sim e^{-\beta v}$, where v is the valency of the welded lattice and excitations have unit mass. A weld breaks then at a rate $r_1 \sim L^2 e^{-\beta v}$, where we include an L^2 term to account for the size of the welded surface. High-energy processes such as creation on a weld are exponentially suppressed with inverse temperature compared with processes that reduce the energy of the system and we might thus expect the Hamiltonian to reverse the effects of thermal errors. Indeed, this is ultimately why the two-dimensional Ising model performs as an effective memory in its ordered phase.

In spite of having favorable energetics that are analogous to those of the Ising model, at low temperatures we expect to observe processes that break the weld at a rate much quicker than r_1, which are due to the large volume of qubits within each code block involved in a weld. These processes are summarized in Fig. 18. To analyze the low-energy processes we examine the surface of a weld. A single excitation created within a code block cannot pass the weld without incurring a high-energy cost. However, if multiple excitations from different blocks meet at the same point along the weld, the excitations can pass through at a low-energy cost. We make use of the ideal gas equation $PV = nRT$ to show that we can expect these low-energy processes to occur most commonly in the low-temperature limit.

We model the excitations that have occurred in a single block of volume $V = L^3$ as an ideal gas of point excitations of
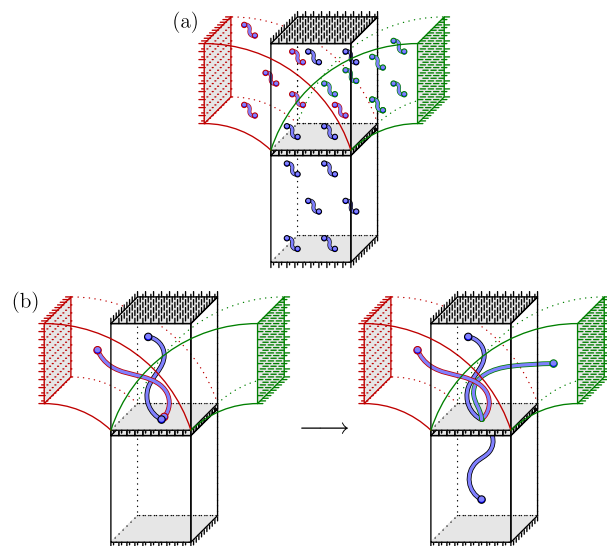

FIG. 18. The excitations of the welded toric code. (a) The blocks that make up the welded code contain an excitation gas of density $\rho \sim e^{-\beta}$ at thermalization. (b) A zero-energy process; if $v/2$ excitations of different blocks meet at a common point on a v-valent weld, a weld can be broken at no energy cost.

density $e^{-\beta}$, as shown in Fig. 18(a), which is achieved quickly as the model approaches equilibrium. Using the fact that the number of point excitations in a given block is $n = Ve^{-\beta}$ and that $T = 1/\beta$, the excitation pressure on the boundaries of a given block follows from the ideal gas equation

$$P \sim e^{-\beta}/\beta, \qquad (77)$$

where we take the gas constant R as unity. We use the pressure to estimate the rate at which excitations fall on a particular point on the boundary to find the rate at which multiple excitations from different blocks of the weld meet at a common point, as shown in Fig. 18(b). If $v/2$ point particles from $v/2$ different blocks that share a welded face meet at a common point, a weld can be overcome at no energy cost. This event occurs at a rate $r_2 \sim L^2 P^{v/2}$, where the L^2 term comes from the size of the welded face. We compare this rate with r_1 to find

$$\frac{r_1}{r_2} \sim \beta^{v/2} e^{-\beta v/2}, \qquad (78)$$

which vanishes in the limit $\beta \to \infty$. This shows that low-energy processes are the dominant processes in the low-temperature limit and we thus argue that the energy barrier will be ineffective in this regime. We obtain the same conclusion by modeling the free motion of excitations using other physically reasonable dynamics.

By consideration of low-energy processes and heuristic calculations we have argued that we cannot easily predict the thermal behavior of the welded model by simple understanding of its nontrivial energy barrier or by application of Arrhenius's law. As such, it is interesting to understand the thermal behavior of the model at intermediate size and temperature regimes. Indeed, while the present manuscript was under peer review, work by Siva and Yoshida (2016) emerged indicating that the welded code will demonstrate superexponential coherence-time scaling with inverse temperature through the study of finite temperature topological order. More generally, a careful study of the dynamics of the model may allow us to derive new no-go theorems for finite-temperature stability that rely on a clearer understanding of entropic mechanisms that decohere encoded quantum information. Recent work following this direction was conducted by Temme (2014) by consideration of the case of commuting Pauli-Hamiltonians.

D. Fractal product codes

We finally remark on the new result of Brell (2016), where *fractal product codes* are introduced. In Brell (2016), the proposed family of models are mapped onto a classical model known to have a finite-temperature phase transition to argue that the model will be stable below a finite critical temperature.

The model is found using the formalism of homological product codes (Bravyi and Hastings, 2013; Freedman and Hastings, 2014). The homological product is an operation that combines pairs of codes to find new codes that in general are locally embeddable only in a larger number of spatial dimensions than their composite parts. The homological product of two two-dimensional toric codes for instance returns the four-dimensional toric code.

The model presented by Brell (2016) is the homological product of two two-dimensional toric codes defined on the Sierpiński carpet graph. This gives a code that resembles the four-dimensional toric code defined on a fractal-like sublattice of the four-dimensional hypercubic lattice. It is conjectured that the choice of graph enables a local embedding of the product code in three dimensions. To demonstrate the stability of the model, it is shown that the model can be mapped onto the partition function of the product code onto an Ising model defined on a Sierpiński carpet, a model which has been rigorously proven to have a finite-temperature phase transition (Shinoda, 2002; Vezzani, 2003; Campari and Cassi, 2010).

The discovery of fractal product codes has opened a new avenue of research, and as such they have raised many new questions. For instance, it is yet to be shown that the model can be efficiently decoded. Moreover, the model has an extensive ground-state degeneracy. This means we cannot easily apply known results to prove that it is perturbatively stable. Indeed, it is not even clear what the geometry of such a code might look like in three spatial dimensions if it is indeed embeddable. Further study of this model may lead to exciting insight toward stable quantum memories that can be realized in the laboratory.

VIII. OTHER PROTECTION MECHANISMS

Following the restrictive no-go theorems described in Sec. IV that forbid nontrivial energy barriers for large classes of two-dimensional commuting systems, it is interesting to consider other mechanisms that inhibit the long-range propagation of errors. Such a study is well motivated as protection mechanisms that do not require a macroscopic energy barrier may be compatible with experimentally amenable two-dimensional models. The purpose of this section is to discuss other proposed mechanisms for the preservation of quantum information that do not rely on macroscopic energy barriers.

In Sec. VIII.A we discuss a model introduced by Brown, Al-Shimary, and Pachos (2014) that is designed to exploit entropic effects to suppress thermal errors from developing. We go on to discuss known limitations of entropic protection in Sec. VIII.B. Finally, in Sec. VIII.C we briefly discuss mechanisms to protect quantum information where the dominant noise source is coherent in nature.

A. Entropically suppressed thermal errors

It is interesting to ask if it is possible to protect quantum information from thermal errors by optimizing the entropic term of the free energy in Eq. (32) to increase the coherence time of a system (Landon-Cardinal and Poulin, 2013), particularly in low-dimensional systems, where the energy barrier between orthogonal ground states is necessarily constant. In this section we describe a two-dimensional model (Brown, Al-Shimary, and Pachos, 2014) that relies exclusively on entropic effects to protect encoded quantum information at finite temperature. We see by consideration of the dynamics of the system that the propagation of the commonly occurring

excitations is suppressed. Here we give a qualitative picture of the mechanism that gives rise to entropic behavior by consideration of the anyonic excitation spectrum of the model. For a technical description of the underlying Hamiltonian, we refer the interested reader to Brown, Al-Shimary, and Pachos (2014).

The entropically protected model makes use of the fusion space of a generalized toric code model defined on a lattice of $L \times L$ vertices with N-level spins on the edges of the lattice. For the reader familiar with quantum double models (Kitaev, 2003), we are considering the quantum double of the group $\mathbb{Z}_N$.

The relevant feature of the generalized toric code model that we discuss here is its anyonic excitation structure. Like the toric code, the generalized model has two types of anyons, electric charges, and magnetic fluxes. However, in the generalized model they carry integer charges $1 \leq k \leq N-1$. We label excitations e_k and m_k. The relevant fusion rules for the model are

$$e_k \times e_l = e_{k \oplus l} \quad \text{and} \quad m_k \times m_l = m_{k \oplus l}, \qquad (79)$$

where $\oplus$ denotes addition modulo N. It follows from these fusion rules that the antiparticles of e_k and m_k excitations are $\overline{e_k} = e_{N-k}$ and $\overline{m_k} = m_{N-k}$. Henceforth, we restrict the discussion to only e_k particles. Because of the symmetries of the model an equivalent discussion holds for m_k particles.

1. Thermal dynamics

As with the case of the toric code, excitations of the generalized model are free to propagate long distances across the lattice at no energy cost. This introduces large errors to the underlying physical lattice and thus rapidly decoheres information encoded in the ground state. Unlike the toric code, the excitations of the generalized model have a splitting structure, where excitations e_k can divide into two spatially separated excitations e_j and e_l provided the charge is conserved, i.e., $j \oplus l = k$; see Eq. (79). We write this splitting process as $e_k \to e_j \times e_l$. This additional structure follows immediately from the fusion rules (79). Errors incident to the lattice can achieve high-energy configurations of many excitations due to the splitting structure of the model. We adapt the generalized toric code to exploit this splitting structure.

To encourage splitting processes to occur when coupled to the thermal environment, we write a Hamiltonian that assigns different masses to excitations of different charges. We choose the masses of the model such that it is energetically favorable for a subset of excitations to split. At this point we consider the explicit case for $N = 5$. We set the Hamiltonian such that masses M_k of particles e_k are

$$2M_1 = 2M_4 \leq M_2 = M_3. \qquad (80)$$

With this setup, it is energetically favorable for the decay processes $e_2 \to e_1 \times e_1$ and $e_3 \to e_4 \times e_4$ to occur.

With the described setup at moderately low temperatures, and given that the model is initialized in the ground state, the most common process that we expect will decohere the information encoded in the ground space is the creation of

an $e_1 \times e_4$ pair that will subsequently propagate rapidly across the lattice. The innovation of Brown, Al-Shimary, and Pachos (2014) is the introduction of defect lines that entropically inhibit the long-range low-energy propagation of excitations by encouraging high-energy splitting processes. Defect lines are studied in generality in Kitaev and Kong (2012) and Barkeshli *et al.* (2014). Loosely speaking, in the general theory of defect lines, Hamiltonian terms are modified along a defect line such that when anyonic excitations cross the defect line, they modify their particle type according to some mapping.

The entropically protected model uses defect lines that modify the charge of crossing excitations. We show two defect lines lying on the lattice in red in Fig. 19(a). Importantly, a defect line maps e_k excitations crossing the line in the negative direction onto $e_{k \oplus k}$ excitations, where the addition is carried out modulo N. Conversely m_k excitations multiply their charge to become $m_{k \oplus k}$ excitations if they cross the defect lines in the opposite direction. The inverse operation occurs if excitations move over the defect line in the reversed direction. Importantly, the defect lines are designed such that if the commonly occurring low-mass e_1 and e_4 excitations cross defect lines in either direction their charge will always be modified to those of high-mass excitations.

We consider the long-range propagation of excitations in a moderately low-temperature regime, where low-mass excitations are generated sparsely across the lattice. In Fig. 19(a) an e_1 excitation is shown in the bottom-left corner of the lattice. Because of the choice of masses (80), commonly occurring
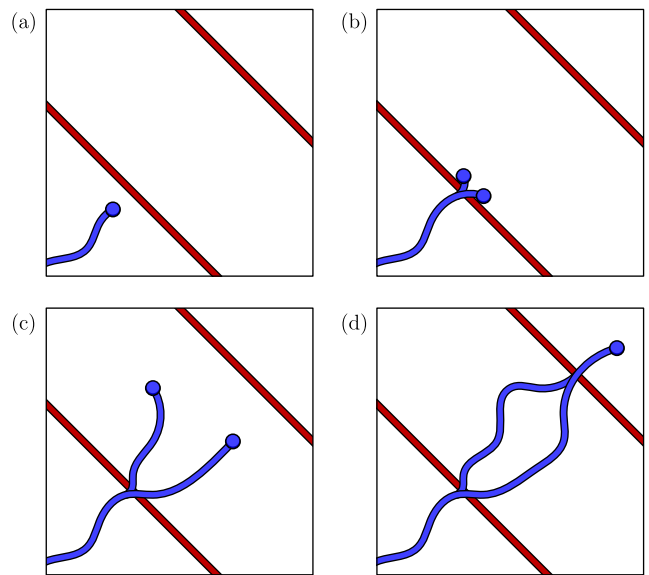

FIG. 19. Entropically suppressed excitations. (a) A single e_1 excitation is marked by a blue circle. Defect lines are drawn in red along diagonal lines of the lattice. (b) The e_1 particle propagates across the defect line to become an e_2 particle. This process is energetically suppressed by the choice of Hamiltonian. (c) A common process for the entropically protected model is for high-mass excitations to decay into pairs of low-mass excitations. The low-mass excitations are confined between the defect lines. (d) In the limit of very low temperatures, the lowest energy process for pairs of low-mass excitations to pass a defect line is by recombination. This limits the entropic effects we describe at low temperatures.

low-mass excitations are energetically suppressed from moving right across a defect line. The energy penalty reduces the rate of decoherence. The energetically suppressed process is shown in Fig. 19(b).

Once the excitation crosses the defect line, it can continue to propagate across the lattice and over the next defect line as an e_2 particle with no energy penalty. However, we can configure the model such that this process is highly improbable. Because of the energetics of the model, it is favorable for the decay process $e_2 \rightarrow e_1 \times e_1$ to occur. This process is shown in Fig. 19(c). Following the decay process, the two e_1 excitations are confined between its enclosing defect lines, as they are subject to an energy penalty to propagate beyond a defect. Given the freedom to choose the defect line separation, we can optimize the system to make this process highly probable, thus commonly suppressing low-energy diffusion of errors. In what follows we discuss the numerics demonstrating the entropically favorable process that we have described thus far.

2. Numerical simulations

Brown, Al-Shimary, and Pachos (2014) simulated the entropically protected model in a thermal environment to show that the typical excitations increase in energy as they diffuse in some appropriate parameter regime. The model is set up with a square grid of defect lines with separations alternating between one and two lattice units where the excitation masses $M_1 = M_4 = 0.38$ and $M_2 = M_3 = 1$ are taken. These values best embellish the entropic effects with numerically tractable system sizes.

The thermalization data obtained by Brown, Al-Shimary, and Pachos (2014) are compared to the partial self-correction hypothesis introduced in Sec. VII.A in the temperature regime $6 \leq \beta \leq 9$. By comparison to Eqs. (73) and (74), the values

$$\Delta_{EP} \sim 0.5 \quad \text{and} \quad \kappa_{EP} \sim 0.2 \qquad (81)$$

are obtained for the entropically protected model. The positive κ_{EP} value is indicative of error dynamics with excitation masses that grow with the total size of the error. We observe that the entropic error suppression shown here is significantly weaker than those found in three-dimensional partially self-correcting models. This is reflected by comparing the obtained data (81) with the data found for the cubic code (76). This can be explained by the fact that entropic protection relies on probabilistic effects for energetic suppression, unlike the cubic code model where an energy barrier is inherent in the system. We further remark that unlike the cubic code model these effects are limited to the regime $\beta \leq 9$, and as such do not satisfy the conditions we require of a quantum memory. This is because at very low temperatures the thermal environment will find low-energy paths to propagate excitations such as that shown in Fig. 19(d). This low-temperature behavior is reminiscent of that which occurs in the welded code, as discussed in Sec. VII.C.1. In the following section we discuss known fundamental limitations on entropic protection.

B. On entropic protection

It is interesting that the entropic protection of the model introduced in the previous section is effective only above a certain moderately high temperature. Following the introduction of the entropically protected model, significant work has been conducted to learn the limitations of error suppression by entropic effects.

Temme (2014) and Temme and Kastoryano (2015) showed for commuting Pauli-Hamiltonians that a macroscopic energy barrier is necessary to achieve a coherence time that diverges with system size. This work was extended by Kómár, Landon-Cardinal, and Temme (2016) to show the necessity of an energy barrier for a more general class of models, namely, Abelian quantum double models (Kitaev, 2003), which include the entropically protected model discussed in Sec. VIII.A. It is thus clear that improvements in coherence time that are achieved by the introduction of defect lines are strictly finite in nature and will never lead to diverging memory time, or super-Arrhenius-law coherence-time scaling in the $\beta \rightarrow \infty$ limit.

Despite the negative results toward self-correction via entropic effects, as two-dimensional quantum memories will almost certainly be more experimentally amenable than their higher-dimensional counterparts, it may still be worthwhile improving low-dimensional memories with entropic effects. As such, it may be interesting to optimize the parameters of entropically protected memories, such as that we reviewed in the previous section, to boost their coherence times.

We finally remark that the known no-go theorems for entropic protection are applicable to Abelian quantum double models. It may be interesting to consider other mechanisms for entropic protection in noncommuting Hamiltonians, or even with commuting models that give rise to non-Abelian anyon theories (Kitaev, 2003; Levin and Wen, 2005). Another interesting proposal for a two-dimensional memory was given by Bardyn and Karzig (2015). They considered coupling the toric code Hamiltonian to driven-dissipative ancilla systems (Pastawski, Clemente, and Cirac, 2011) to inhibit the propagation of excitations. It is the role of the dissipative systems to reduce the energy of the interaction terms of the Hamiltonian to introduce potential minima to the system to increase the likelihood of diffusing excitations retracing their steps which will thus reverse incident errors. It was shown in Bardyn and Karzig (2015) that coupling the toric code to a dissipative ancilla system gives rise to superexponential coherence-time scaling with inverse temperature.

C. Coherent noise suppression

The majority of this review has been concerned with finding systems that protect quantum information encoded in many-body systems from incoherent sources of noise, namely, where the system of interest is coupled to a thermal bath. However, we might consider some implementation of a system where finite-temperature effects are negligible, and where coherent sources of noise are dominant. Specifically, we consider here weak local perturbations such as external magnetic fields. Protection against coherent sources of noise will become increasingly relevant as the temperature of the

system becomes very low. In such a regime we might consider designing systems where we sacrifice the self-correction and instead focus on experimental amenability. Here we discuss mechanisms that have been considered to defend low-dimensional Hamiltonian systems from various forms of coherent noise.

As discussed in Sec. II.F, perturbative stability at zero temperature is a natural property of topologically ordered systems. It is this inherent stability that first motivated their use for quantum computation (Kitaev, 2003). In such systems, quantum information can be stored within locally indistinguishable degenerate ground states, which are separated from excited states by a finite gap. Perturbations will cause some energy splitting between ground states and can cause the gap to close to a degree. However, it has been shown that these effects will be suppressed with system size for perturbations whose strength is below some finite threshold (Bravyi, Hastings, and Michalakis, 2010). Such topological systems, including the simple toric code, therefore provide a stable memory at zero temperature.

Similar properties also arise in so-called symmetry protected topological phases (Chen, Gu, and Wen, 2011; Else, Bartlett, and Doherty, 2012; Bonderson and Nayak, 2013). These offer protection against perturbations that respect nontrivial symmetries of the model. As compensation for this, they are typically more experimentally accessible than topological models that do not rely on a symmetry.

A well-known example is an open chain for spinless fermions with a superconducting pairing term (Kogut, 1979; Kitaev, 2001). For a suitable choice of Hamiltonian parameters, this system can be brought to a phase supporting topological superconductivity (Kitaev, 2001) with a perturbatively stable degenerate ground state and Majorana fermion zero-energy modes localized at opposite ends of the chain.

Quantum information encoded within this subspace spanned by these modes (Kitaev, 2001; Bravyi, 2006) will remain stable as long as the fermionic parity of the wire is conserved, which relies on only Cooper pairs being exchanged with the environment. Exchange of single unpaired fermions can cause fatal decoherence (Rainis and Loss, 2012). The system must therefore be engineered to ensure that such symmetry-violating processes are rare.

So far, low-dimensional symmetry protected phases have not been shown to offer self-correction against thermal noise. In fact, certain models have been found to decohere in a way that is not efficiently suppressed by reducing temperature (Campbell, 2015). In such cases, it has been argued that active error correction is the only hope of preserving the symmetry protected quantum memory at finite temperature (Pedrocchi and DiVincenzo, 2015).

This discussion on the suppression of coherent noise assumes that information can be encoded directly into the ground-state manifold. However, in practice we must expect that preparing the state will be noisy such that we may prepare an excited state with some mobile quasiparticle excitation on the lattice. It is thus problematic that the coherent noise can propagate these excitations and decohere stored information very quickly.

One mechanism studied to suppress this is disorder-assisted protection. Pastawski *et al.* (2010), Stark *et al.* (2011), Wootton and Pachos (2011a), Bravyi and König (2012), and Röthlisberger *et al.* (2012) showed that randomizing the Hamiltonian interaction strengths inhibits the coherent propagation of excitations across the lattice. This is attributed to the Anderson localization effect (Anderson, 1958), where it is understood that randomized Hamiltonian interactions lead to "friction" in the motion of excitation dynamics due to quasiparticles becoming trapped in small potential minima of the random energy landscape. It is tempting then to believe that randomness might then be useful as a resource for error suppression. However, Anderson localization is not a well-understood principle. Remarkably, Bravyi and König (2012) showed that pseudorandom potentials outperform truly randomly chosen potentials.

Some of the randomness studied by Röthlisberger *et al.* (2012) was of the underlying lattice structure in the two-dimensional toric code model, rather than the coupling strengths. This demonstrates the important effect that the chosen lattice will have on coherence time, as expanded upon in Al-Shimary, Wootton, and Pachos (2013). Although this focuses primarily on thermal noise, its main purpose is to optimize the lattice in the case in which noise is biased toward certain kinds of errors, specifically, either bit flip or dephasing noise. It is identified that reducing the connectivity of the lattice geometry which embeds the toric code will reduce the rate of one type of noise from decohering the lattice. However, due to the dual structure of the toric code lattice, changing the lattice geometry to protect against one type of noise will detrimentally affect the performance of the lattice against the other type of noise. This approach therefore presents a tradeoff between protection against bit flip and dephasing noise introduced by a general thermal noise model. Therefore these effects provide a constant improvement in coherence in the presence of an asymmetric noise model that may be present in a realistic experimental setting (Douçot, Feigel'man, and Ioffe, 2003).

IX. BEYOND COMMUTING HAMILTONIANS: SUBSYSTEM CODES

In this section we discuss progress toward the study of a class of noncommuting models that has developed over the last decade, namely, the subsystem codes. This framework provides a natural extension to the stabilizer formalism that we have relied upon throughout this review. Among the subsystem code literature are three-dimensional models that are conjectured to be self-correcting quantum memories.

Subsystem codes were introduced by Kribs, Laflamme, and Poulin (2005) and Kribs *et al.* (2006). The language of subsystems was introduced to find a unifying language for decoherence free subspaces (Palma, Suominen, and Ekert, 1996; Duan and Guo, 1997; Zanardi and Rasetti, 1997; Lidar, Chuang, and Whaley, 1998; Beige *et al.*, 2000; Pachos and Beige, 2004) and noiseless subsystems (Knill, Laflamme, and Viola, 2000; Zanardi, 2000; Kempe *et al.*, 2001). Initially coined operator quantum error correction, subsystem codes principally encode logical information in a quantum error-correcting code that is embedded in a subsystem of a larger

Hilbert space. In a subsystem code the remainder of the Hilbert space of the total system is referred to as the *gauge subsystem*.

A general language for subsystem codes is provided by Poulin (2005). A subsystem code is uniquely defined by its gauge group $\mathcal{G}_n$, a subgroup of the Pauli group of n qubits. Given a gauge group, a stabilizer code is defined on the subsystem acted upon by the centralizer of the gauge group $\mathcal{N}(\mathcal{G}_n)$, i.e., operators that commute with elements of $\mathcal{G}_n$. The stabilizers of the code are also members of the gauge group $\mathcal{S} = \mathcal{N}(\mathcal{G}_n) \cap \mathcal{G}_n$, whereas logical operators $\mathcal{L} = \mathcal{N}(\mathcal{G}_n) \backslash \mathcal{G}_n$ commute with but are not themselves members of the gauge group. It is easily checked that for the special case in which $\mathcal{G}_n$ is Abelian, we recover the stabilizer formalism where $\mathcal{G}_n = \mathcal{S}$ up to phases.

Throughout this section we give examples which show that this innocuous abstraction of the stabilizer formalism behaves qualitatively differently from the stabilizer error-correcting model. First, by definition, the gauge subsystem can take any state. It can, for instance, become arbitrarily mixed due to incident noise or otherwise, and encoded quantum information will remain robust. One might also consider using the gauge subsystem for other practical purposes. Herrera-Martì, Li, and Kwek (2014) made frequent measurements on the subsystem surface code presented by Bravyi *et al.* (2013) to suppress the coherent diffusion of excitations at zero temperature. It is also shown that the subsystem structure of certain codes can be used to realize universal quantum computation (Paetznick and Reichardt, 2013), or to design advantageous error-correcting protocols (Bombin, 2014) which we discuss in Sec. IX.C.

We additionally remark that subsystem codes may have some differences in their practical realization compared with their stabilizer counterparts which must be taken into account if one is to build a quantum memory based on a subsystem code. In particular, certain local subsystem codes give rise to high-weight stabilizer measurements that become increasingly error prone in the large system-size limit. We urge the interested reader to find a comprehensive review of this topic in Terhal (2015) and references therein to better understand the practical difficulties in realizing certain subsystem error-correcting codes. Conversely, some subsystem codes (Bombin, 2010b; Bravyi *et al.*, 2013) are specifically designed to reduce the weight of syndrome measurements such that they are less demanding to implement from an experimental perspective.

Little is known about the fundamental features or thermal characteristics of subsystem codes, although progress has been made in this area by Chesi *et al.* (2010) where bounds on their relaxation times are obtained. Much of the work in this area has been a constructive search for models that we might expect to give rise to favorable properties for self-correction. Of particular interest is the three-dimensional Bacon-Shor code (Bacon, 2006). This model has drawn significant attention to subsystem codes as it is conjectured to behave as a self-correcting memory. This argument is made as its Hamiltonian has key features in common with thermally stable classical models, namely, the Ising model, discussed in Sec. V.A. We introduce the three-dimensional Bacon-Shor

code by first reviewing the two-dimensional Bacon-Shor code, a simple example of a subsystem code that has some qualitatively different features from local stabilizer codes. We conclude our discussion of the two-dimensional Bacon-Shor code by considering its generalizations and other subsystem codes that are shown to surpass constraints physically imposed on general commuting models.

Subsystem codes are of further interest when considered in the context of topological order. For the reader familiar with topologically ordered lattice models, we remark that Kitaev's honeycomb model (Kitaev, 2006) provides an illustrative example of a subsystem code (Suchara, Bravyi, and Terhal, 2011); the model is studied by considering loop degrees of freedom that commute with its noncommuting parent Hamiltonian. The honeycomb model falls into a broader subclass of subsystem codes, known as *topological subsystem codes* (Bombin, 2010b; Suchara, Bravyi, and Terhal, 2011). Recently a three-dimensional generalization of topological subsystem codes, namely, the gauge color code model, has been conjectured to be self-correcting. We conclude this section with an overview of this model and its features that have led to this conjecture. We also discuss other interesting results that have arisen by consideration of the gauge color code as a self-correcting quantum memory.

A. The Bacon-Shor code

The Bacon-Shor code (Bacon, 2006), otherwise known as the quantum compass model (Kugel and Khomskii, 1982; Dorier, Becca, and Mila, 2005), provides a nontrivial example of a subsystem code that demonstrates physical features not accessible with local stabilizer codes. For an extensive review of compass models see Nussinov (2013). In the two-dimensional model, we observe that its local gauge generators give rise to nonlocal stabilizer generators. Moreover, we see how the nontrivial gauge group affects the error-correction procedure. We conclude this section by reviewing results that show the fundamental differences between subsystem codes and commuting models.

The two-dimensional Bacon-Shor code is defined on an $L \times L$ square lattice with qubits on its vertices, as shown in Fig. 20. The gauge group $\mathcal{G}_n$ where $n = L^2$ is generated by two types of operator. It has two-body nearest-neighbor interactions A_x and B_x associated with lattice sites $x = (j, k)$, where $A_{j,k} = X_{j,k} X_{j+1,k}$ are aligned along the horizontal direction and $B_{j,k} = Z_{j,k} Z_{j,k+1}$ operators are aligned in the vertical direction, as shown in Fig. 20(a) in green and blue, respectively. We write the corresponding Hamiltonian of this model as

$$H_{2\text{DBS}} = -\sum_x (A_x + B_x). \tag{82}$$

In general, A_x and B_x operators do not commute. The stabilizers of this code $\mathcal{S} = \mathcal{N}(\mathcal{G}_n) \cap \mathcal{G}_n$ are nonlocal operators of the gauge group generated by products of A_x and B_x operators. They are $S_j^X = \prod_k A_{j,k}$ and $S_k^Z = \prod_j B_{j,k}$. The stabilizers S_j^X and S_k^Z are supported on bands, two vertices wide, such as those shown in green and blue in Figs. 20(b) and 20(c).

The horizontal bands that support S_k^Z stabilizers have four common qubits with the support of the vertical bands supporting S_j^X stabilizers. As such, it is easily checked that all stabilizers commute. The model encodes a single logical qubit. Its logical operators $\mathcal{L} = \mathcal{N}(\mathcal{G}_n)\backslash\mathcal{G}_n$ are $\bar{X} = \prod_k X_{j,k}$ for fixed j and $\bar{Z} = \prod_j Z_{j,k}$ for fixed k. The support of $\bar{X}$ and $\bar{Z}$ are shown in red and yellow in Fig. 20. They cross at a single point by Fig. 20(d).

The two-dimensional Bacon-Shor code shows how subsystem codes respond differently to noise. We consider a noise model that introduces Pauli-X errors. An equivalent discussion holds for Pauli-Z errors where the lattice is rotated by $\pi/2$. Remarkably, for any given row of the lattice, we need only to correct the *parity* of errors. A pair of Pauli-X errors along a horizontal strip are elements of $\mathcal{G}_n$ and thus commute with all $\mathcal{S}$ and $\mathcal{L}$. Errors of this type therefore do not affect information encoded in the subsystem code. As in the case for stabilizer codes, we make stabilizer measurements to identify the parity of errors between pairs of strips, and we correct the parity of a given row using a single-qubit Pauli-X operation along that row.

Importantly, as pointed out in the original paper (Bacon, 2006), it was argued persuasively that the gap of Hamiltonian (82) vanishes in the thermodynamic limit. This is shown in Dorier, Becca, and Mila (2005) using extensive numerical methods in an extended region of phase space. The model is therefore only likely to perform well as a quantum memory in some limited regime where the system size is relatively small. Nevertheless, the example shows that a clever choice of gauge enables us to ignore large classes of errors that affect only the gauge subsystem. Further, the Bacon-Shor code provides an

example of a local subsystem code model that gives rise to a nonlocal stabilizer code, thus indicating a fundamental difference between local stabilizer and subsystem codes.

Indeed, the study of local subsystem codes that give rise to nonlocal stabilizer codes has been significantly extended due to the recent work by Bacon *et al.* (2015). In this work they show a very general scheme where one can find a local subsystem code that gives rise to a stabilizer code, that is not necessarily local, given the quantum circuit that measures the stabilizer generators of the stabilizer code. They use their formalism to find codes with favorable code distance scaling that saturates known bounds for commuting projector codes. As remarked by Bacon *et al.* (2015), perhaps one could consider using their formalism to construct local subsystem codes that correspond to nonlocal stabilizer models with favorable properties at finite temperature to find stable quantum memories. Certainly, it is shown that the presented construction can give rise to models with no stringlike logical operators, even if one considers the more general class of dressed logical operators, i.e., logical operators with nontrivial action on the gauge subsystem of the model.

We finally remark on an extension of the Bacon-Shor code due to Bravyi *et al.* (2013), where he shows that the fundamental storage capacity of a two-dimensional subsystem code can surpass the storage capacity of commuting models (Bravyi, Poulin, and Terhal, 2010). In the commuting case, it is known that $kd^2 \leq \mathcal{O}(n)$, where k is the number of qubits of the code, and d is the minimum weight of the lowest weight logical operator of the code. Bravyi shows that in the case of subsystem codes one can obtain scaling like $kd \leq \mathcal{O}(n)$ with a local gauge group, and that randomly chosen codes saturate this bound asymptotically in the limit of large n. Unfortunately, this bound cannot be saturated for codes with constant k. Indeed it is also known from Bravyi and Terhal (2009) and Haah and Preskill (2012) that the distance of two-dimensional subsystem codes must satisfy $d^2 \leq \mathcal{O}(n)$. Results such as these provide further intrigue and motivation toward the study of subsystem codes, as we clearly see that these noncommuting codes are capable of supporting encoding properties provably unattainable by their commuting counterparts.

B. The three-dimensional Bacon-Shor code

Subsystem codes have attracted a lot of interest from the point of view of thermal stability due to the three-dimensional Bacon-Shor code. Bacon (2006) conjectured that the model gives rise to self-correcting behavior in a thermal setting.

The model is defined on an $L \times L \times L$ lattice for odd L with sites labeled by $x = (j, k, l)$. The corresponding Hamiltonian is given by

$$H_{3\mathrm{DBS}} = -\sum_x (A_x + B_x + C_x + D_x), \qquad (83)$$

where the interaction terms are two-body nearest-neighbor Pauli operators $A_{j,k,l} = X_{j,k,l} X_{j+1,k,l}$, $B_{j,k,l} = X_{j,k,l} X_{j,k+1,l}$, $C_{j,k,l} = Z_{j,k,l} Z_{j,k+1,l}$, and $D_{j,k,l} = Z_{j,k,l} Z_{j,k,l+1}$ which generate the gauge group. The logical operators of the code are two-dimensional planelike operators $\bar{X} = \prod_{k,l} X_{1,k,l}$ and $\bar{Z} = \prod_{j,k} Z_{j,k,1}$ that anticommute due to the condition that L is odd.

FIG. 20. The Bacon-Shor code in two dimensions. Qubits are represented by white circles. (a) Example of gauge generators. The gauge group $\mathcal{G}_n$ is generated by nearest-neighbor two-body operators including the horizontal A_x operators and the vertical nearest-neighbor B_x operators which are shown, respectively, in green and blue. The shaded regions by (b) and (c) show the support of stabilizers of the Bacon-Shor code. (d) The support of $\bar{X}$ and $\bar{Z}$ is shown by the narrow, red shaded vertical strip and the narrow, yellow shaded horizontal strip, respectively. The logical operators intersect at a single qubit marked by dark orange shading near (d).

The stabilizer generators are two-dimensional planelike operators that are two vertices wide, such that $S_j^X = \prod_{k,l} X_{j,k,l} X_{j+1,k,l}$ and $S_l^Z = \prod_{j,k} Z_{j,k,l} Z_{j,k,l+1}$.

The Hamiltonian of the three-dimensional Bacon-Shor code has the gauge terms of many intersecting two-dimensional Ising models. It has two-body Pauli-X–type interactions along one two-dimensional plane, and Pauli-Z–type interactions along an orthogonal plane. We show the construction in Fig. 21. Bacon and Casaccino (2006) showed that in general one can design generalized Bacon-Shor codes, where one takes an arbitrary classical code input to find its gauge group. Its overlying stabilizer code then inherits the properties of the input classical code. The intuition of the three-dimensional Bacon-Shor code is that it inherits the stability properties of its input code, the two-dimensional Ising model, as discussed in Sec. V.A. Bacon (2006) considered a simple Pauli noise model and mean-field arguments to show that the model may demonstrate a macroscopic power-law energy barrier. The question of stability of the three-dimensional Bacon-Shor code remains an open problem.

The challenge of interrogating this conjecture, as pointed out by Bacon (2006), is that the noise model considered is not representative of all possible channels under which encoded information can decohere. It is possible that a finite-temperature environment may find a low-energy path to introduce errors to the encoded logical information. Indeed, no exact diagonalization of Hamiltonian (83) has been found and without one it is difficult to make strong statements about the model.
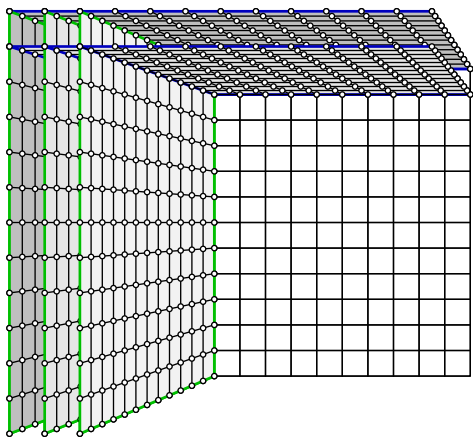
FIG. 21. The gauge group of the three-dimensional Bacon-Shor Hamiltonian. The three-dimensional Bacon-Shor code is supported on qubits arranged on cubic lattice. The gauge group consists of two-body nearest-neighbor terms, where different planes of the lattice support different types of gauge operators. Horizontal planes such as those we outline in blue support two-body Pauli-X gauge operators between their qubits. Similarly, the vertical planes outlined in green that protrude out of the page support two-body Pauli-Z gauge terms along its edges. The full model is described by an array of L of these planes arranged in parallel along their respective directions. The locations of the full set of planes are marked by black lines at the back of the figure. Indeed, the gauge group of the three-dimensional Bacon-Shor code can be regarded as intersecting copies of the two-dimensional Ising model supported on each of these planes.

Certainly, flaws have been identified with the model that must be overcome to satisfy the required conditions of a stable memory. First, the model provably has no error-correction procedure in the thermodynamic limit (Pastawski *et al.*, 2010). Further, given that we are relying on the model to inherit the stability properties of the two-dimensional Ising model, we may also expect it to inherit its perturbative instability, described in Richards *et al.* (1995), Cirillo and Lebowitz (1998), Grinstein (2004), and Pastawski *et al.* (2010). One might support this assertion by extrapolating results regarding the gap from the well-studied two-dimensional case (Dorier, Becca, and Mila, 2005). If the model is indeed gapless, like the two-dimensional code, we may encounter issues with the perturbative stability of the three-dimensional Bacon-Shor code as we increase the size of the system toward the thermodynamic limit.

C. The gauge color code

We conclude this section with a discussion of the *gauge color code* (Bombin, 2015), the smallest realization of which was first discovered by Paetznick and Reichardt (2013) by consideration of Reed-Muller codes. The model represents a three-dimensional topological subsystem code. Topological subsystem codes are discussed in two dimensions in Bombin (2010b), Brell *et al.* (2011), Suchara, Bravyi, and Terhal (2011), Sarvepalli and Brown (2012), Bravyi *et al.* (2013), and Kubica and Beverland (2015). The model is of particular interest as it is conjectured to give rise to finite-temperature stability (Bombin, 2015). For a comprehensible overview of the gauge color code see Kubica and Beverland (2015) and Watson *et al.* (2015).

We briefly elaborate on the fault-tolerant computational properties of the color codes. Notably the color code models are favorable for their implementation of transversal gates. The two-dimensional color code has a transversal implementation of the Clifford gate set (Bombin and Martin-Delagado, 2006) which, importantly here, includes the Hadamard gate. In the two-dimensional model, the feature that enables its implementation is the self-duality of its stabilizers. Explicitly, a self-dual stabilizer group is such that for every stabilizer $S^X = \prod_{j \in \mathcal{T}} X_j$ supported on a subset of qubits $\mathcal{T}$, there exists also the stabilizer $S^Z = \prod_{j \in \mathcal{T}} Z_j$.

The three-dimensional stabilizer color code (Bombin and Martin-Delagado, 2007b) gives rise to a transversal implementation of the controlled-not gate and the $\pi/8$ gate. Supplemented by the Hadamard gate, the gate set of the three-dimensional color code is capable of universal fault-tolerant quantum computation. However, the three-dimensional stabilizer color code is not self-dual, and as such does not support a universal transversal gate set.

Bombin (2015) showed that a universal gate set can be realized with a subsystem generalization of the three-dimensional color code. Importantly, the stabilizer group of the three-dimensional color code contains a self-dual subset of stabilizers that are capable of successfully identifying an arbitrary set of errors below a certain weight. Given a suitable choice of gauge generators, we can restrict the gauge color code stabilizers to its self-dual subset, such that we

have a code with a transversal implementation of the Hadamard gate.

The proposal of Bombin represents an explicit lattice realization of gauge fixing, introduced by Paetznick and Reichardt (2013). Gauge fixing avoids resource costly methods of achieving a universal gate set using, for instance, magic state distillation (Bravyi and Kitaev, 2005). Instead, gauge fixing enables us to fault tolerantly move information between different codes by changing the gauge operators of a given subsystem code. Specifically, one can "promote" certain elements of the gauge group of a code to elements of the stabilizer group by imposing that the code states of the code take particular eigenstates of gauge operators, thus changing, or fixing, the code. The development of gauge fixing has shown that we can move between different error-correcting codes that collectively support a universal gate set. Gauge fixing is shown explicitly for the Reed-Muller codes in Anderson, Duclos-Cianci, and Poulin (2014). The idea of gauge fixing extends ideas presented by Knill, Laflamme, and Zurek (1996) where fault-tolerant gates that do not preserve the code space are suggested.

Here we consider the gauge color code for a candidate self-correcting quantum memory. A lattice suitable to describe the code is shown in Fig. 22. Another appropriate lattice geometry is given by Kim (2011); see also Brown, Nickerson, and Browne (2016). Importantly, the lattice is four valent and four colorable, i.e., the lattice is such that we can assign to each cell one of four colors in such a way that no two neighboring cells are of the same color. It follows that three-dimensional lattices that satisfy these properties have faces that contain an even number of vertices (Kubica and Beverland, 2015). The generators of the gauge group are associated with the faces f of cells of the lattice. Specifically, for every face of a cell, we have gauge generators $A_f = \prod_{v \in f} X_v$ and $B_f = \prod_{v \in f} Z_v$, where we use shorthand $v \in f$ to mean the vertices adjacent to face f. On the lattice shown, gauge generators are either four- or six-body terms. We therefore have the Hamiltonian

$$H_{\mathrm{GCC}} = -\sum_f (A_f + B_f). \tag{84}$$

The stabilizers are then associated with the cells c of the lattice, such that $A_c = \prod_{v \in c} X_v$ and $B_c = \prod_{v \in c} Z_v$, where $v \in c$ are qubits associated with vertices adjacent to cell c.

The conjecture due to Bombin (2015) is based on the structure of the A_f and B_f Hamiltonian terms. Specifically, if one considered the simpler Hamiltonian $H^X = -\sum_f B_f$ subject only to the Pauli-X error channel, then the excitations of the model are exclusively akin to Peierls contours, as discussed in Sec. V. An equivalent argument holds for Pauli-Z noise for Hamiltonian $H^Z = -\sum_f A_f$, and as such one can argue that there may be a macroscopic energy barrier for arbitrary local quantum noise channels. Of course, given the difficulty in diagonalizing Hamiltonian (84) this argument is not rigorous.

When defined on a tetrahedron, the logical operators that commute with the gauge generators are two dimensional. We show the smallest tetrahedron that embeds a gauge color code in Fig. 23. Both $\bar{X}$ and $\bar{Z}$ are supported on one face of the tetrahedral lattice. Their support is shown in blue (dark shaded) on the example lattice in Fig. 23(a). The geometry of the color code ensures that the support of a face of the tetrahedron remains odd for any size of lattice.

While it is a promising feature with respect to self-correction that the logical operators are two dimensional, it is not clear that such logical operators are sufficient for self-correction in subsystem codes. Indeed, the support of logical operators of subsystem codes can be reduced by multiplication by gauge generators. Such a logical operator will not necessarily commute with the gauge group, but still commutes with the stabilizer group, as required. A logical operator of this type is known as a dressed logical operator.

In contrast to the models proposed by Bacon *et al.* (2015) mentioned in Sec. IX.A, the gauge color code has one-dimensional dressed logical operators. We show the support of the dressed logical operators of the gauge color code in Fig. 23(b). We can conclude little from discovering this operator, but we make this point only to illustrate some of the additional complexity involved in studying subsystem codes. Indeed, due to the nontrivial commutation relations of the dressed operator with respect to the interaction terms of Hamiltonian (84), it is not clear that such an operator can be achieved at a constant energy cost under a local noise channel. As such, it is not known if the discovery of a low-dimensional

FIG. 22. The lattice of the gauge color code. Qubits reside on the vertices of a four-valent lattice. The lattice is also four colorable. This means we can assign each of the three-dimensional cells of the lattice one of four colors such that no two cells of the same color touch. The cells are colored dark blue, red, light green, and yellow to show the four colorability of the cells.
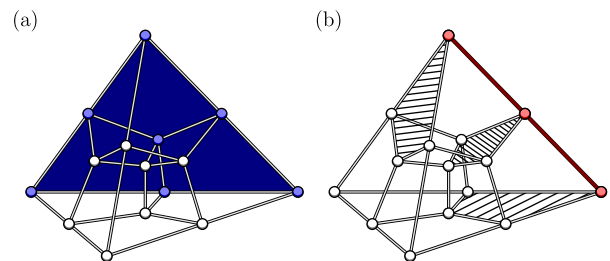

FIG. 23. A 15-qubit tetrahedral lattice that supports the gauge color code. (a) The two-dimensional logical operator of the gauge color code, marked in blue with shaded vertices at the back of the lattice. (b) The support of the dressed logical operator is shaded red. The faces which support Hamiltonian terms that anticommute with the dressed logical operator are shaded.

dressed logical operator rules out the possibility of self-correction in the gauge color code model. We additionally remark that, like the three-dimensional Bacon-Shor code, it is not even obvious that the gauge color code is gapped. To this end, the thermal stability of the gauge color code remains an open problem.

We finally mention further progress in the study of the gauge color code with respect to its error-correcting capabilities. Bombin (2014) showed that the gauge color code model has favorable properties for decoding when measured using unreliable laboratory equipment. He argued that we need to measure each face operator of the gauge color code only once to obtain reliable fault-tolerant syndrome information. This phenomenon, coined single-shot error correction, has been demonstrated numerically by Brown, Nickerson, and Browne (2016). This differs from the well-studied case of two-dimensional stabilizer models where syndrome information is read out using an unreliable measurement apparatus. Known schemes for fault-tolerant error correction require that each stabilizer must be measured a macroscopic number of times to read out logical information reliably (Dennis *et al.*, 2002; Wang, Harrington, and Preskill, 2003). Bombin remarked that the special error-correcting properties he demonstrated are generic to known self-correcting stabilizer models (Bombin, 2014). Certainly, this is an exciting advance in active error correction that has been derived from the study of a candidate self-correcting model.

X. DISCUSSION AND OUTLOOK

Among the major discoveries that led to scalable classical information processing was the discovery of a transistor. The physics of these little solid-state devices ensure the reliable storage and robust processing of classical information that is easily scaled. The discovery of an experimentally amenable stable quantum memory presents a monumental hurdle, which, if overcome, will be invaluable for the discovery of fault-tolerant quantum information processing (Bacon, 2006). In the present review we have given an overview of the analytical and numerical tools we use to approach the study of many-body quantum systems at finite temperature. We have examined the no-go theorems that have been discovered in this field, and we have presented many new physical models with certain properties suitable for the passive protection of quantum information. We conclude by highlighting some of the forthcoming challenges we face toward discovering and developing a quantum memory.

This review has highlighted many open problems in this field. We have discussed models that arguably present favorable properties to make them resilient in a thermal environment. Such models include the ferromagnet-coupled toric code, the welded toric code, embeddable-fractal product codes, the three-dimensional Bacon-Shor code, and the gauge color code. These conjectures need to be interrogated by numerical experiments or by rigorous proofs. To study these models we need to develop both analytical and numerical techniques in both condensed-matter physics and fundamental statistical mechanics. In addition to this, it needs to be checked that these models have other features that are required of a passive quantum memory, such as efficient decoding algorithms and perturbative stability.

A noteworthy theme that has occurred frequently throughout this review is that many of the candidate models for finite-temperature stability are composed of simpler systems. For instance, the work in Sec. VI shows that we can introduce anyonic interactions with a local model by coupling the toric code to either a system of bosons or a Heisenberg ferromagnet. Similarly, the nontranslationally invariant models in Secs. VII.C and VII.D are constructed by combining simpler topologically ordered models using either welding or by taking the homological product of many codes. Indeed, the entropically protected model presented in Sec. VIII.A can be regarded as a patchwork lattice of many different, albeit equivalent, topological phases. Moreover, the Bacon-Shor codes enable the combination of favorable classical models using the underlying gauge structure of subsystem codes. To develop new models one might consider adapting these composition tools further to construct new hybrid Hamiltonians with features we require of a quantum memory.

Finally, we conclude by emphasizing that all the promising theoretical models that are proposed must be developed until they are sympathetic to the engineering constraints of the laboratory. Certain proposals for a finite-temperature quantum memory are more amenable to experimentalists than others, for instance, the ferromagnetic coupled toric code was designed with physical media in mind. Moreover, subsystem codes offer an avenue to simplify the architectural challenges of building high-weight stabilizer models. However, in spite of a few exceptions, experimentally amenable quantum memories that can be realized using existing technology is an avenue of research which thus far remains untrodden. We consider the example of the celebrated cubic code model, a well-studied model whose favorable partial self-correcting properties have been analyzed and numerically verified. We must work now to develop models such as this one into a form that an experimentalist can prepare in the laboratory. Undoubtedly, such an achievement will be directly incorporated in the design of fault-tolerant quantum technologies of the future.

APPENDIX A: DECODERS

Throughout this review we make use of decoding algorithms for numerical analysis. Given an encoded quantum state that has been subject to noise, a decoder takes classical syndrome information, i.e., all the outcomes of stabilizer measurements, and returns a correction operation that returns the state to the code space (Dennis *et al.*, 2002). Provided the errors occur at a suitably low density, the decoder will successfully find a correction operator that will recover the initially encoded state with probability that grows with the size of the system. Here we briefly discuss decoding, and a particular decoding routine, namely, the clustering decoder. This present discussion makes use of the stabilizer formalism described in Sec. II.

Even in the case that we find a self-correcting quantum memory, a decoding step will still be required to correct for errors caused by small thermal fluctuations when information is read from the system. To illustrate this, we briefly consider the two-dimensional Ising model, as discussed in Sec. V.A. We encode classical bits in the twofold degenerate ground space, and we read out by measuring the magnetization, the average spin value of all the spins of the system. The ground states for the model have magnetization ± 1. At finite temperature in the limit of large system sizes we cannot expect all the spins to be aligned. Instead it is suitable to take the sign of the magnetization measurement to read out the memory. This measurement corresponds to taking a majority vote over all the encoded spins of the lattice. Measuring the encoded ferromagnet in this way accounts for small thermal fluctuations that take the spin configuration out of the ground space. For robust quantum information storage we require a decoding algorithm to deal with small errors incident to a code during readout.

As we cannot measure the state of individual physical qubits of a code, accounting for the errors during the readout of a quantum code is not as straightforward as the classical case we have described. Instead, for the quantum case, we perform stabilizer measurements to learn the errors that are incident to a code. The stabilizer measurements perform the task of collapsing the incident noise onto an error E that is an element of the Pauli group and provides syndrome information we can use to attempt to determine E. It is the task of the decoder to predict the error E of the Pauli group and return a correction operator C such that CE acts trivially on the encoded state.

Many approaches to decoding have been studied with tradeoffs between speed, performance, and versatility. Decoders have been designed that make use of minimum-weight perfect matching (Dennis *et al.*, 2002), renormalization group techniques (Duclos-Cianci and Poulin, 2010, 2013), and Monte Carlo methods (Wootton and Loss, 2012; Hutter, Wootton, and Loss, 2014). Moreover, the study of decoding has a foundation in the study of glassy statistical mechanical models (Dennis *et al.*, 2002; Wang, Harrington, and Preskill, 2003; Bombin *et al.*, 2012; Andrist, Wootton, and Katzgraber, 2014).

Here we describe the clustering decoder which is commonly used throughout this review. The clustering decoder was introduced by Harrington (2004) and developed by

Bravyi and Haah (2011a, 2013) specifically for the study of the cubic code at finite temperature, as discussed in Sec. VII.B. The clustering decoder was further refined in Anwar *et al.* (2014) and Hutter, Loss, and Wootton (2014). This simple algorithm is particularly suitable for the present work as it can be adapted for any translationally invariant local stabilizer code (Bravyi and Haah, 2011a).

To find a correction operator for the most likely error configuration, the clustering decoder will implicitly make use of the locality constraint of commuting Pauli-Hamiltonian models. In addition to this, we assume that a low-weight correction operator will approximate the inverse of the most probable error that has occurred in the limit of a low error rate.

Here we sketch the clustering algorithm routine. A rigorous explanation of the implementation of the decoder can be found in Bravyi and Haah (2011a), Anwar *et al.* (2014), and Hutter, Loss, and Wootton (2014). In Fig. 24(a) we show a series errors, marked in white, and the syndrome that corresponds to the error. Violated stabilizers, i.e., stabilizers that return -1 outcomes, are marked by black points for some local code. We consider this example to demonstrate the clustering decoder.

To find a low-weight correction we find a set of small boxes that enclose all of the error syndromes. We search for a set of boxes that contains a correction operator that is consistent with the violated stabilizers in each box. The algorithm begins by putting all the violated stabilizers in individual boxes of unit size, as shown in Fig. 24(b). The algorithm then proceeds by incrementally increasing the size of the boxes. This is achieved by combining pairs of boxes that lie within some small radius of one another. The routine continues until the boxes are large enough to contain a correction operator that corrects for all the violated stabilizers contained within the box.
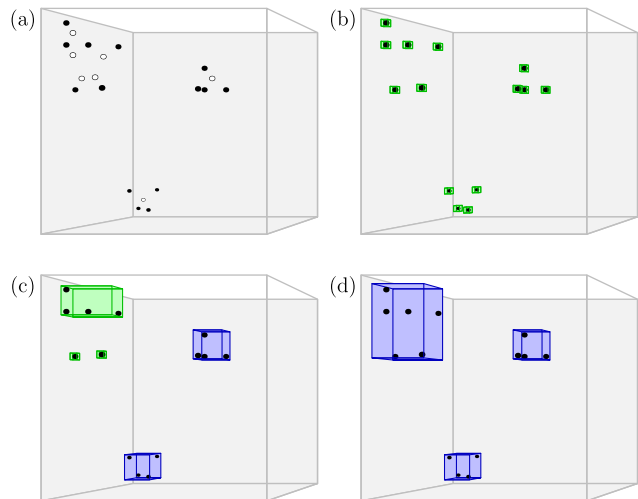

FIG. 24.   The clustering decoding algorithm. (a) Unknown errors, marked in white on a three-dimensional lattice, are identified by syndrome measurements, marked in black. (b) The measured syndromes are initially contained in unit boxes. The unknown errors are not shown. (c) Box sizes increase to contain other nearby syndromes within a small fixed radius of the existing boxes. Boxes that contain a correction operator are colored dark blue (right, bottom); they are otherwise colored light green. (d) The search increases the box size to find boxes large enough to contain correction operators for all the syndromes on the lattice.

For the initially chosen boxes, where all of the boxes are of unit size, we cannot find a local correction operator that is consistent with the single violated stabilizer that is contained within each box. We increase the size of the boxes by checking within a fixed radius r of each of the violated stabilizers of each box. In the event that another violated stabilizer contained in a different box is found within a distance r of the violated stabilizer from which we are searching, their two respective boxes are combined giving a single larger box. Once it is confirmed that no pairs of disjoint boxes contain any violated stabilizers that lie within distance r of one another, we check to find a correction operator that is consistent with all of the violated stabilizers in each box.

If a box contains a correction operator that is consistent with all of its violated stabilizers, the box is considered neutral, and the violated stabilizers within the box are no longer considered in later searches of the routine. In the case that all the boxes are neutral, the algorithm terminates and the correction operator of each of the boxes is returned. If some boxes remain unneutralized, r is increased and the algorithm is repeated for the violated stabilizers that are contained in boxes that are not yet neutral.

We continue to follow the example syndrome in Fig. 24. In Fig. 24(c) we show the new boxes obtained after searching within a radius of $r = 1$ of each of the violated stabilizers. After this search, not all the boxes contain a local correction operator. Two boxes, colored blue in Fig. 24(c), are neutralized. Violated stabilizers in the blue boxes are no longer considered in the algorithm. Three boxes found at radius $r = 1$, marked in green, do not contain a correction operator consistent with their respective violated stabilizers. To find a correction operator consistent with the violated stabilizers in the green boxes, the algorithm increases its box search radius once again to $r = 2$ for all the remaining violated stabilizers. At $r = 2$, all boxes contain a correction operator that is consistent with all of the violated stabilizers. Figure 24(d) shows the boxes that together contain a correction operator consistent with the syndromes of the error configuration. The decoder will return the correction operator contained within the blue boxes to return the encoded information to the ground state, enabling the readout of the encoded quantum information.

We remark that this algorithm is suitable for any translationally invariant local stabilizer code. The description of the algorithm we give makes no reference to the underlying code. We need only assume that the code is local, such that violated stabilizers can be interpreted as lying within a fixed radius of an incident error. Moreover, Bravyi and Haah (2011a) showed that for transitionally invariant codes one can efficiently determine if a box contains a correction operator consistent with its enclosed violated stabilizers.

We can evaluate decoder performance by obtaining a *threshold* with respect to an identical and independently distributed noise model. This noise model is where each qubit suffers an error with probability p. The threshold value p_{th} is the value below which the logical error rate of a quantum error-correcting code decreases as we increase the size of the system to the thermodynamic limit. In this review we use cluster decoding for the toric code model in Sec. III.F and the cubic
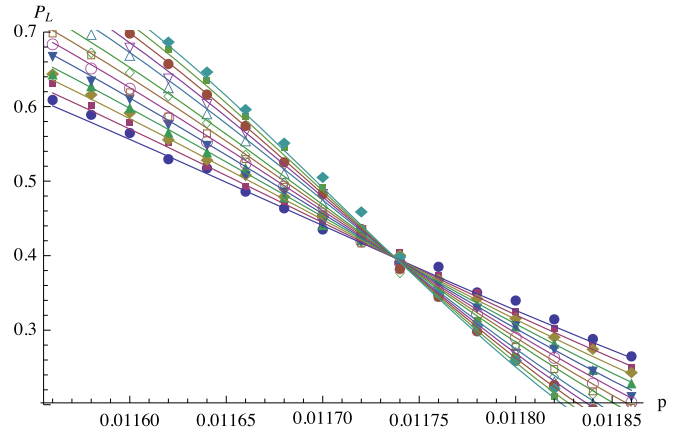

FIG. 25. Threshold calculation using the clustering decoder for the cubic code model. We study an independent and identically distributed bit-flip noise model using system sizes between $L = 101$ and 201 with Monte Carlo samples. We find the crossing at $p_{th} = 1.17\%$.

code in Sec. VII.B. The toric code threshold is already found to be $\sim 8.3\%$, shown in Anwar *et al.* (2014). The threshold for the cubic code using the clustering algorithm has not been published to the best of our knowledge, although good estimates are given in Bravyi and Haah (2011a). We estimate a threshold $p_{th} \sim 1.17\%$ for the cubic code using a bit-flip noise model. We show the threshold data in Fig. 25.

APPENDIX B: SIMULATING THE TORIC CODE AT FINITE TEMPERATURE

Here we present additional numerical results for the finite-temperature behavior of the toric code, supporting the arguments we make in Sec. III.F.

For system sizes L that are small compared to the natural scale imposed by the finite-temperature dynamics, decoherence is typically the result of a single pair of anyons. In Sec. III.F we predicted scaling with L and inverse temperature $\beta = 1/T$ of three independent elements of the coherence time, which we called $\Pi(L, \beta)$, τ_c, and τ_m. We isolate each of these terms and estimate them numerically.

We first investigate $\Pi(L, \beta)$. This term quantifies the probability that after a pair of anyons is created they do not annihilate by mutually fusing together before reaching a significant enough distance to cause logical errors $L/2$. To evaluate this function we alter the standard simulation scheme; rather than beginning in a ground state we initialize the system with a single pair of anyons present on the lattice, where the initialized anyons are separated by a single lattice spacing. We additionally set the rate of creation equal to zero such that no additional pairs of anyons are created. Indeed, we model only the random motion of a single pair of anyons walking across the lattice. We evolve the system until either the separation of the anyons reaches the Euclidean distance $L/2$ or the pair meet at a common point on the lattice and annihilate. The quantity Π is the fraction of samples that reach separation $L/2$ rather than annihilating. We estimate Π by sampling over 10^4 trials. The results we obtain are shown in Fig. 26. We find the fitting
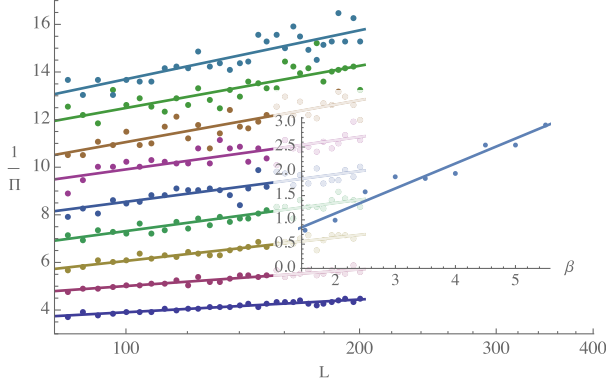
FIG. 26. The reciprocal of the probability showing that a pair of toric code excitations reach critical separation $L/2$. Values of Π are obtained by averaging over 10 000 simulations. They are plotted against L over a range of temperatures from $\beta = 1$ (bottom line) to $\beta = 6$ (top line). The linear fittings show that $1/\Pi$ grows linearly with $\ln(L/2)$, with a gradient that increases with β. The inset shows the gradients found with the fittings shown in the main plot displayed as a function of β.

$$1/\Pi = (0.108 + 0.513\beta)\ln(L/2). \tag{B1}$$

This is in good agreement with the scaling we hypothesized in Eq. (43), where constant $A \approx 5$.

We next numerically study how τ_c and τ_m scale with L and β. Time τ_c is the typical time it takes for a pair of anyons to be created that will cause a logical error, and τ_m is the average time it takes for a pair of anyons to achieve separation $L/2$ after they have been created. To find these values we simulate the thermal evolution of the system prepared in a ground state. We attempt to decode the lattice after each simulated event of the thermal evolution, and the simulation is stopped once the decoder fails. Every time a pair of anyons is created its creation time is recorded; this information is discarded if the pair annihilates. Estimates of the creation time scale τ_c and diffusion time τ_m are obtained by averaging over 10^3 simulation runs. These data are plotted in Figs. 7 and 27, respectively. The data we obtain scale as we expect with fittings given in Eqs. (42) and (44), where we find numerical fits $\tau_m \simeq 0.028\beta L^2$ and

$$\tau_c = 0.150 \frac{e^{1.99\beta}}{L^{2.01}} \frac{1}{\Pi}, \tag{B2}$$

where we take the function Π that we evaluated numerically, Eq. (B1).

For large systems, decoherence results due to the interaction of many thermally created anyons. In Sec. III.F we describe a model of these dynamics that gives an estimate, Eq. (48), of the coherence time τ_{large}. We test some of the assumptions of this model and compare the predicted coherence time to numerical values.

We simulate the system evolving in the large size regime. The decoder runs after every event are introduced to the system by the thermal evolution. At the earliest time the decoder fails we stop the simulation and record the time elapsed during the simulation. Results are obtained by
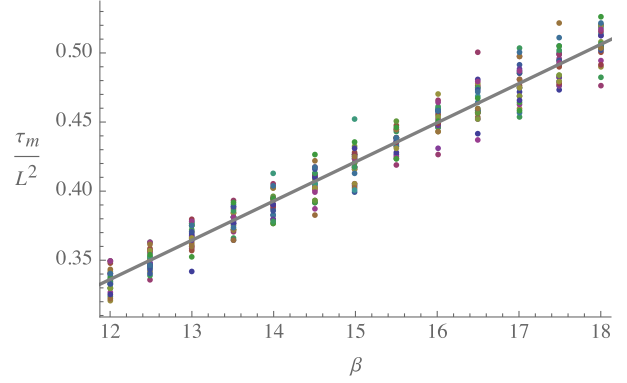

FIG. 27. Contribution to the coherence time from anyon motion in the small system-size limit. Times τ_m are obtained by averaging over 1000 simulations. They are plotted against β for a range of system sizes $L = 50, 52, 54, \ldots, 100$. The values of τ_m are divided by L^2 to show data points that are independent of the size of the system, thus validating the L^2 factor we derive on the right-hand side of Eq. (42). The linear fit displayed uses the average values of the fit parameters obtained from the different system sizes $\tau_m \sim 0.028\beta L^2$.

averaging over 10^3 simulations runs. We find average anyon numbers $\langle N \rangle$ which confirm that the anyon density at the time the decoder fails scales like $\rho \sim e^{-\beta\Delta}$. These numbers are also seen to satisfy $\langle N \rangle \gg 1$, indicating that our data are taken for sufficiently large systems.

In Sec. III.F we argued that for large systems, in contrast to the smaller case, the important length scale is $\Lambda \sim e^{\beta\Delta/2}$ as opposed to system size. This value is the typical separation between creation events and thus corresponds to the average distance each anyon must diffuse to cause the decoder to fail.

To test that the described motion is the main decoherence mechanism for large systems we study the distance between anyon pairs at the time the decoder fails. During the simulation each pair of anyons is given a unique mark to
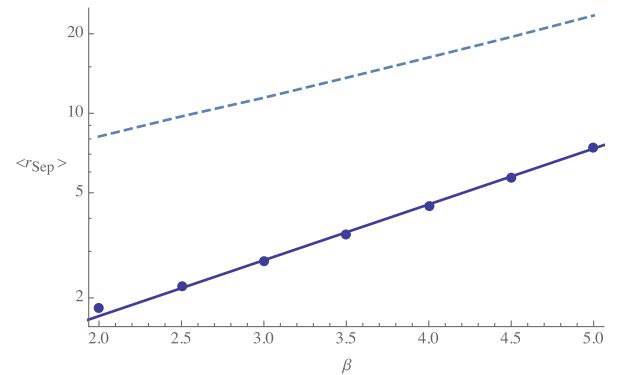

FIG. 28. Average separation between paired anyons at the time the decoder fails in the large system-size limit. The separation is plotted as a function of β for fixed system size $L = 200$ where each data point is obtained by averaging over 1000 simulations. The solid line is a linear fit to the data with gradient ~ 0.49 on logarithmic axes, consistent with a scaling of $\langle r_{\text{sep}} \rangle \sim e^{\beta\Delta/2}$. The dashed line is the maximum pair separation, averaged over the simulations, which is seen to be much smaller than $L/2$.
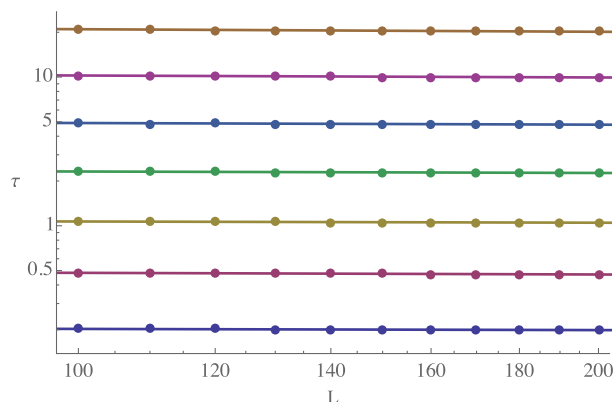
FIG. 29. The coherence time of the toric code is independent of system size at high temperature. The quantity τ is shown on a log scale against L for different temperatures $\beta = 2, 2.5, \ldots, 5$ with $\beta = 2$ the bottom line and $\beta = 5$ the top line. Data points are found by taking an average over 1000 simulations. Linear fits show very small negative gradients ranging from -0.03 for the lowest value of β down to -0.06 for the highest β. Indeed, the zero gradient is well within the 95% confidence interval of each fit.

indicate their pair created partner. At the end of the simulation we measure the Euclidean distance between each marked pair. If two anyons from two separate pairs fuse, the remaining two unpaired anyons on the lattice are marked as members of the same extended pair.

As one might expect, we find that extended pairs created from a fusion will typically achieve a greater separation than pairs that are initialized by creation from vacuum. However, its effect is small with respect to the average data. Our numerical results show that the average separation between all anyon pairs grows like $e^{\beta \Delta / 2}$, as predicted. In addition we find that the maximum separation between any pair is always much less than $L/2$. The scalings of the average and maximum separations are shown in Fig. 28. These observations confirm that the decoherence results from the average motion of anyons in local regions.

The scaling of coherence time with β is shown in Fig. 8. We see that the data reproduce the exponential dependence on $\beta\Delta$ predicted in Eq. (48). Another prediction of the model is that τ_{large} is independent of system size. Figure 29 plots the numerical values of coherence times against L showing that the data are consistent with this hypothesis.

REFERENCES

Aharonov, D., and M. Ben-Or, 1997, in *STOC '97 Proceedings of the twenty-ninth annual ACM symposium on Theory of computing* (ACM, New York), p. 176.

Alicki, R., 2012, Open Syst. Inf. Dyn. **19**, 1250016.

Alicki, R., and M. Fannes, 2009, Phys. Rev. A **79**, 012316.

Alicki, R., M. Fannes, and M. Horodecki, 2007, J. Phys. A **40**, 6451.

Alicki, R., M. Fannes, and M. Horodecki, 2009, J. Phys. A **42**, 065303.

Alicki, R., and M. Horodecki, 2006, arXiv:quant-ph/0603260.

Alicki, R., M. Horodecki, P. Horodecki, and R. Horodecki, 2010, Open Syst. Inf. Dyn. **17**, 1.

Alicki, R., and K. Lendi, 2007, *Quantum Dynamical Semigroups and Applications* (Springer, New York).

Aliferis, P., and A. W. Cross, 2007, Phys. Rev. Lett. **98**, 220502.

Aliferis, P., D. Gottesman, and J. Preskill, 2006, Quantum Inf. Comput. **6**, 97.

Al-Shimary, A., J. R. Wootton, and J. K. Pachos, 2013, New J. Phys. **15**, 025027.

Anderson, J. T., G. Duclos-Cianci, and D. Poulin, 2014, Phys. Rev. Lett. **113**, 080501.

Anderson, J. T., and T. Jochym-O'Connor, 2014, arXiv:1409.8320.

Anderson, P. W., 1958, Phys. Rev. **109**, 1492.

Andrist, R. S., J. R. Wootton, and H. G. Katzgraber, 2014, arXiv:1406.5974.

Anwar, H., B. J. Brown, E. T. Campbell, and D. E. Browne, 2014, New J. Phys. **16**, 063038.

Bacon, D., 2006, Phys. Rev. A **73**, 012340.

Bacon, D., and A. Casaccino, 2006, arXiv:quant-ph/0610088.

Bacon, D., S. T. Flammia, A. W. Harrow, and J. Shi, 2015, in *SOTC '15 Proceedings of the forty-seventh annual ACM on Theory of computing* (ACM, New York), p. 327.

Bardyn, C.-E., and T. Karzig, 2015, arXiv:1512.04528.

Barends, R., *et al.*, 2014, Nature (London) **508**, 500.

Barkeshli, M., P. Bonderson, M. Cheng, and Z. Wang, 2014, arXiv:1410.4540.

Barkeshli, M., C.-M. Jian, and X.-L. Qi, 2013, Phys. Rev. B **87**, 045130.

Becker, D., T. Tanamoto, A. Hutter, and F. L. P. D. Loss, 2013, Phys. Rev. A **87**, 042340.

Beige, A., D. Braun, B. Tregenna, and P. L. Knight, 2000, Phys. Rev. Lett. **85**, 1762.

Bombin, H., 2010a, Phys. Rev. Lett. **105**, 030403.

Bombin, H., 2010b, Phys. Rev. A **81**, 032301.

Bombin, H., 2014, arXiv:1404.5504.

Bombin, H., 2015, New J. Phys. **17**, 083002.

Bombin, H., R. S. Andrist, M. Ohzeki, H. G. Katzgraber, and M. A. Martin-Delgado, 2012, Phys. Rev. X **2**, 021004.

Bombin, H., R. W. Chhajlany, M. Horodecki, and M. A. Martin-Delgado, 2013, New J. Phys. **15**, 055023.

Bombin, H., and M. A. Martin-Delagado, 2006, Phys. Rev. Lett. **97**, 180501.

Bombin, H., and M. A. Martin-Delagado, 2007a, Phys. Rev. B **75**, 075103.

Bombin, H., and M. A. Martin-Delagado, 2007b, Phys. Rev. Lett. **98**, 160502.

Bombin, H., and M. A. Martin-Delgado, 2009, J. Phys. A **42**, 095302.

Bonati, C., 2014, Eur. J. Phys. **35**, 035002.

Bonderson, P., and C. Nayak, 2013, Phys. Rev. B **87**, 195451.

Bortz, A. B., M. H. Kalos, and J. L. Lebowitz, 1975, J. Comput. Phys. **17**, 10.

Bouwmeester, D., *et al.*, 1999, Phys. Rev. Lett. **82**, 1345.

Bravyi, S., 2006, Phys. Rev. A **73**, 042313.

Bravyi, S., 2011, Phys. Rev. A **83**, 012320.

Bravyi, S., D. P. DiVincenzo, and D. Loss, 2011, Ann. Phys. (Amsterdam) **326**, 2793.

Bravyi, S., D. P. DiVincenzo, D. Loss, and B. M. Terhal, 2008, Phys. Rev. Lett. **101**, 070503.

Bravyi, S., G. Duclos-Cianci, D. Poulin, and M. Suchara, 2013, Quantum Inf. Comput. **13**, 0963.

Bravyi, S., and J. Haah, 2011a, arXiv:1112.3252.

Bravyi, S., and J. Haah, 2011b, Phys. Rev. Lett. **107**, 150504.

Bravyi, S., and J. Haah, 2013, Phys. Rev. Lett. **111**, 200501.

Bravyi, S., and M. B. Hastings, 2013, arXiv:1311.0885.

Bravyi, S., M. B. Hastings, and S. Michalakis, 2010, J. Math. Phys. (N.Y.) **51**, 093512.

Bravyi, S., and A. Kitaev, 2005, Phys. Rev. A **71**, 022316.

Bravyi, S., and R. König, 2012, Commun. Math. Phys. **316**, 641.

Bravyi, S., and R. König, 2013, Phys. Rev. Lett. **110**, 170503.

Bravyi, S., B. Leemhuis, and B. M. Terhal, 2011, Ann. Phys. (Amsterdam) **326**, 839.

Bravyi, S., D. Poulin, and B. Terhal, 2010, Phys. Rev. Lett. **104**, 050503.

Bravyi, S., and B. Terhal, 2009, New J. Phys. **11**, 043029.

Bravyi, S. B., and A. Y. Kitaev, 1998, arXiv:quant-ph/9811052.

Brell, C. G., 2016, New J. Phys. **18**, 013050.

Brell, C. G., S. Burton, G. Dauphinais, S. T. Flammia, and D. Poulin, 2014, Phys. Rev. X **4**, 031058.

Brell, C. G., S. T. Flammia, S. D. Bartlett, and A. C. Doherty, 2011, New J. Phys. **13**, 053039.

Brennen, G. K., and J. K. Pachos, 2008, Proc. R. Soc. A **464**, 1.

Brown, B. J., A. Al-Shimary, and J. K. Pachos, 2014, Phys. Rev. Lett. **112**, 120503.

Brown, B. J., N. H. Nickerson, and D. E. Browne, 2016, Nat. Commun. **7**, 12302.

Burton, S., C. G. Brell, and S. T. Flammia, 2015, arXiv:1506.03815.

Calderia, A. O., and A. J. Leggett, 1981, Phys. Rev. Lett. **46**, 211.

Campari, R., and D. Cassi, 2010, Phys. Rev. E **81**, 021108.

Campbell, E. T., 2015, in *10th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2015)*, edited by S. Beigi and R. Koenig (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany), Vol. 44 of Leibniz International Proceedings in Informatics (LIPIcs), pp. 111–126.

Castelnovo, C., and C. Chamon, 2007, Phys. Rev. B **76**, 184442.

Castelnovo, C., and C. Chamon, 2008, Phys. Rev. B **78**, 155120.

Castelnovo, C., and C. Chamon, 2012, Philos. Mag. **92**, 304.

Chamon, C., 2005, Phys. Rev. Lett. **94**, 040402.

Chen, X., Z.-C. Gu, and X.-G. Wen, 2011, Phys. Rev. B **83**, 035107.

Chesi, S., D. Loss, S. Bravyi, and B. M. Terhal, 2010, New J. Phys. **12**, 025013.

Chesi, S., B. Röthlisberger, and D. Loss, 2010, Phys. Rev. A **82**, 022305.

Cirillo, E., and J. Lebowitz, 1998, J. Stat. Phys. **90**, 211.

Córcoles, A. D., E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow, 2015, Nat. Commun. **6**, 6979.

Davies, E. B., 1974, Commun. Math. Phys. **39**, 91.

Day, F. V., and S. D. Barrett, 2012, arXiv:1201.0390.

Dennis, E., A. Kitaev, A. Landahl, and J. Preskill, 2002, J. Math. Phys. (N.Y.) **43**, 4452.

Deutsch, D., 1985, Proc. R. Soc. A **400**, 97.

Devoret, M. H., and R. J. Scoelkopf, 2013, Science **339**, 1169.

DiVincenzo, D. P., and D. Loss, 2005, Phys. Rev. B **71**, 035318.

Dorier, J., F. Becca, and F. Mila, 2005, Phys. Rev. B **72**, 024448.

Douçot, B., M. V. Feigel'man, and L. B. Ioffe, 2003, Phys. Rev. Lett. **90**, 107003.

Duan, L.-M., and G.-C. Guo, 1997, Phys. Rev. Lett. **79**, 1953.

Duclos-Cianci, G., and D. Poulin, 2010, Phys. Rev. Lett. **104**, 050504.

Duclos-Cianci, G., and D. Poulin, 2013, Phys. Rev. A **87**, 062338.

Dusuel, S., M. Kamfor, R. Orús, K. P. Schmidt, and J. Vidal, 2011, Phys. Rev. Lett. **106**, 107203.

Eastin, B., and E. Knill, 2009, Phys. Rev. Lett. **102**, 110502.

Einarsson, T., 1990, Phys. Rev. Lett. **64**, 1995.

Else, D. V., S. D. Bartlett, and A. C. Doherty, 2012, New J. Phys. **14**, 113016.

Farrow, T., and V. Vedral, 2014, arXiv:1406.0659.

Feynman, R. P., 1982, Int. J. Theor. Phys. **21**, 467.

Freedman, M. H., and M. B. Hastings, 2014, Quantum Inf. Comput. **14**, 144.

Freedman, M. H., M. Larsen, and Z. Wang, 2002, Commun. Math. Phys. **227**, 605.

Freeman, C. D., C. M. Herdman, D. J. Gorman, and K. B. Whaley, 2014, Phys. Rev. B **90**, 134302.

Frigerio, A., 1978, Commun. Math. Phys. **63**, 269.

Fuchs, G. D., G. Burkard, P. V. Klimov, and D. D. Awschalom, 2011, Nat. Phys. **7**, 789.

Gardiner, C. W., 1983, *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences* (Springer-Verlag, Berlin).

Garrahan, J. P., and M. E. J. Newman, 2000, Phys. Rev. E **62**, 7670.

Gottesman, D., 2001, Stabilizer Codes and Quantum Error Correction, Ph.D. thesis, California Institute of Technology.

Greenberger, D. M., M. Horne, and A. Zeilinger, 1989, *Bell's Theorem, Quantum Theory and Conceptions of the Universe*, edited by M. Kafatos (Kluwer, Dordrecht), p. 69.

Griffiths, R. B., 1964, Phys. Rev. **136**, A437.

Grinstein, G., 2004, IBM J. Res. Dev. **48**, 5.

Haah, J., 2011, Phys. Rev. A **83**, 042330.

Haah, J., 2013, Commun. Math. Phys. **324**, 351.

Haah, J., 2014, Phys. Rev. B **89**, 075119.

Haah, J., and J. Preskill, 2012, Phys. Rev. A **86**, 032308.

Hamma, A., C. Castelnovo, and C. Chamon, 2009, Phys. Rev. B **79**, 245122.

Hamma, A., R. Ionicioiu, and P. Zanardi, 2005a, Phys. Rev. A **71**, 022315.

Hamma, A., R. Ionicioiu, and P. Zanardi, 2005b, Phys. Lett. A **337**, 22.

Hamma, A., P. Zanardi, and X.-G. Wen, 2005, Phys. Rev. B **72**, 035307.

Harrington, J. W., 2004, Analysis of quantum error-correcting codes: sympletic lattice codes and toric codes, Ph.D. thesis, California Institute of Technology.

Harty, T. P., D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, 2014, arXiv:1403.1524.

Hastings, M. B., 2011, Phys. Rev. Lett. **107**, 210501.

Hastings, M. B., G. H. Watson, and R. G. Melko, 2014, Phys. Rev. Lett. **112**, 070501.

Hastings, M. B., D. Wecker, B. Bauer, and M. Troyer, 2014, arXiv:1403.1439.

Herold, M., E. T. Campbell, J. Eisert, and M. J. Kastoryano, 2014, arXiv:1406.2338.

Herrera-Martì, D. A., Y. Li, and L. C. Kwek, 2014, Quantum Inf. Comput. **14**, 1136.

Huang, K., 1987, *Statistical Mechanics* (John Wiley & Sons, New York).

Hutter, A., and D. Loss, 2014, Phys. Rev. A **89**, 042334.

Hutter, A., D. Loss, and J. R. Wootton, 2014, arXiv:1410.4478.

Hutter, A., F. L. Pedrocchi, J. R. Wootton, and D. Loss, 2014, Phys. Rev. A **90**, 012321.

Hutter, A., J. R. Wootton, and D. Loss, 2014, Phys. Rev. A **89**, 022326.

Hutter, A., J. R. Wootton, B. Röthlisberger, and D. Loss, 2012, Phys. Rev. A **86**, 052340.

Iblisdir, S., D. Pérez-García, M. Aguardo, and J. K. Pachos, 2009, Phys. Rev. B **79**, 134303.

Iblisdir, S., D. Pérez-García, M. Aguardo, and J. K. Pachos, 2010, Nucl. Phys. B **829**, 401.

Ising, E., 1925, Z. Phys. **31**, 253.

Jouzdani, P., E. Novais, I. S. Tupitsyn, and E. R. Mucciolo, 2014, Phys. Rev. A **90**, 042315.

Kapit, E., J. T. Chalker, and S. H. Simon, 2014, arXiv:1408.0959.

Kay, A., 2009, Phys. Rev. Lett. **102**, 070503.

Kay, A., 2011, Phys. Rev. Lett. **107**, 270502.

Kay, A., and R. Colbeck, 2008, arXiv:0810.3557v1.

Kelly, J., *et al.*, 2015, Nature (London) **519**, 66.

Kempe, J., D. Bacon, D. A. Lidar, and K. B. Whaley, 2001, Phys. Rev. A **63**, 042307.

Kim, I. H., 2011, Phys. Rev. A **83**, 052308.

Kim, I. H., 2012, arXiv:1202.0052.

Kim, I. H., and J. Haah, 2016, Phys. Rev. Lett. **116**, 027202.

Kitaev, A., 2006, Ann. Phys. (Amsterdam) **321**, 2.

Kitaev, A., and L. Kong, 2012, Commun. Math. Phys. **313**, 351.

Kitaev, A., and J. Preskill, 2006, Phys. Rev. Lett. **96**, 110404.

Kitaev, A. Y., 1997, Russ. Math. Surv. **52**, 1191.

Kitaev, A. Y., 2001, Phys. Usp. **44**, 131.

Kitaev, A. Y., 2003, Ann. Phys. (Amsterdam) **303**, 2.

Kloeffel, C., and D. Loss, 2013, Annu. Rev. Condens. Matter Phys. **4**, 51.

Knill, E., R. Laflamme, and L. Viola, 2000, Phys. Rev. Lett. **84**, 2525.

Knill, E., R. Laflamme, and W. Zurek, 1996, arXiv:quant-ph/9610011.

Knill, E., R. Laflamme, and W. H. Zurek, 1998, Proc. R. Soc. A **454**, 365.

Kochmański, M., T. Paszkiewicz, and S. Wolski, 2013, Eur. J. Phys. **34**, 1555.

Koenig, R., G. Kuperberg, and B. W. Reichardt, 2010, Ann. Phys. (Amsterdam) **325**, 2707.

Kogut, J., and L. Susskind, 1975, Phys. Rev. D **11**, 395.

Kogut, J. B., 1979, Rev. Mod. Phys. **51**, 659.

Kómár, A., O. Landon-Cardinal, and K. Temme, 2016, arXiv:1601.01324.

Kossakowski, A., 1972, Rep. Math. Phys. **3**, 247.

Kossakowski, A., A. Frigerio, V. Gorini, and M. Verri, 1977, Commun. Math. Phys. **57**, 97.

Kribs, D., R. Laflamme, and D. Poulin, 2005, Phys. Rev. Lett. **94**, 180501.

Kribs, D. W., R. Laflamme, D. Poulin, and M. Lesosky, 2006, Quantum Inf. Comput. **6**, 382.

Kubica, A., and M. E. Beverland, 2015, Phys. Rev. A **91**, 032330.

Kugel, K. I., and D. I. Khomskii, 1982, Sov. Phys. Usp. **25**, 231.

Laidler, K. J., 1972, J. Chem. Educ. **49**, 343.

Landon-Cardinal, O., and D. Poulin, 2013, Phys. Rev. Lett. **110**, 090502.

Landon-Cardinal, O., B. Yoshida, D. Poulin, and J. Preskill, 2015, Phys. Rev. A **91**, 032303.

Lebowitz, J. L., and A. E. Mazel, 1998, J. Stat. Phys. **90**, 1051.

Leggett, A. J., S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg, and W. Zwerger, 1987, Rev. Mod. Phys. **59**, 1.

Leinaas, J. M., and J. Myrheim, 1977, Nuovo Cimento Soc. Ital. Fis. B **37**, 1.

Levin, M., and X.-G. Wen, 2006, Phys. Rev. Lett. **96**, 110405.

Levin, M. A., and X.-G. Wen, 2005, Phys. Rev. B **71**, 045110.

Lidar, D. A., and T. A. Brun, 2013, Eds., *Quantum Error Correction* (Cambridge University Press, Cambridge, England).

Lidar, D. A., I. L. Chuang, and K. B. Whaley, 1998, Phys. Rev. Lett. **81**, 2594.

Lindblad, G., 1976, Commun. Math. Phys. **48**, 119.

Lodyga, J., P. Mazurek, A. Grudka, and M. Horodecki, 2015, Sci. Rep. **5**, 8975.

Loss, D., and D. P. DiVincenzo, 1998, Phys. Rev. A **57**, 120.

Mazáč, D., and A. Hamma, 2012, Ann. Phys. (Amsterdam) **327**, 2096.

McCoy, B. M., and T. T. Wu, 2014, *The Two-Dimensional Ising Model: Second Edition* (Dover Publications, New York).

Michnicki, K. P., 2012, arXiv:1208.3496.

Michnicki, K. P., 2014, Phys. Rev. Lett. **113**, 130501.

Nakahara, M., 2003, *Geometry, Topology and Physics* (Institute of Physics, Bristol/Philadelphia).

Nayak, C., S. H. Simon, A. Stern, M. Freedman, and S. D. Sarma, 2008, Rev. Mod. Phys. **80**, 1083.

Newman, M. E. J., and C. Moore, 1999, Phys. Rev. E **60**, 5068.

Ni, X., O. Buerschaper, and M. V. den Nest, 2014, arXiv:1404.5327.

Nigg, D., M. M. E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt, 2014, arXiv:1403.5426.

Nussinov, Z., 2013, arXiv:1303.5922.

Nussinov, Z., and G. Ortiz, 2008, Phys. Rev. B **77**, 064302.

Nussinov, Z., and G. Ortiz, 2009a, Proc. Natl. Acad. Sci. U.S.A. **106**, 16944.

Nussinov, Z., and G. Ortiz, 2009b, Ann. Phys. (Amsterdam) **324**, 977.

Nussinov, Z., G. Ortiz, and E. Cobanera, 2012, Ann. Phys. (Amsterdam) **327**, 2491.

Onsager, L., 1944, Phys. Rev. **65**, 117.

Pachos, J. K., 2012, *Introduction to Topological Quantum Computation* (Cambridge University Press, Cambridge, England).

Pachos, J. K., and A. Beige, 2004, Phys. Rev. A **69**, 033817.

Paetznick, A., and B. W. Reichardt, 2013, Phys. Rev. Lett. **111**, 090505.

Palma, G. M., K.-A. Suominen, and A. K. Ekert, 1996, Proc. R. Soc. A **452**, 567.

Pastawski, F., L. Clemente, and J. I. Cirac, 2011, Phys. Rev. A **83**, 012304.

Pastawski, F., A. Kay, N. Schuch, and I. Cirac, 2009, Phys. Rev. Lett. **103**, 080501.

Pastawski, F., A. Kay, N. Schuch, and I. Cirac, 2010, Quantum Inf. Comput. **10**, 0580.

Pastawski, F., and B. Yoshida, 2015, Phys. Rev. A **91**, 012305.

Pedrocchi, F. L., S. Chesi, and D. Loss, 2011, Phys. Rev. B **83**, 115415.

Pedrocchi, F. L., and D. P. DiVincenzo, 2015, Phys. Rev. Lett. **115**, 120402.

Pedrocchi, F. L., A. Hutter, J. R. Wootton, and D. Loss, 2013, Phys. Rev. A **88**, 062313.

Peierls, R., 1936, Proc. Cambridge Philos. Soc. **32**, 477.

Poulin, D., 2005, Phys. Rev. Lett. **95**, 230504.

Poulin, D., M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, 2014, arXiv:1406.4920.

Preskill, J., 1998, Proc. R. Soc. A **454**, 385.

Preskill, J., 2004, Lecture Notes on Quantum Computation, Chap. 9 [http://www.theory.caltech.edu/people/preskill/ph219/].

Rainis, D., and D. Loss, 2012, Phys. Rev. B **85**, 174533.

Raussendorf, R., J. Harrington, and K. Goyal, 2006, Ann. Phys. (Amsterdam) **321**, 2242.

Reed, M. D., L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, 2012, Nature (London) **482**, 382.

Richards, L. H., S. W. Sides, M. A. Novotny, and P. A. Rikvold, 1995, J. Magn. Magn. Mater. **150**, 37.

Röthlisberger, B., J. R. Wootton, R. M. Heath, J. K. Pachos, and D. Loss, 2012, Phys. Rev. A **85**, 022313.

Saeedi, K., S. Simmons, J. Z. Salvail, P. Dluhy, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, J. J. L. Morton, and M. L. W. Thewalt, 2013, Science **342**, 830.

Sarvepalli, P., and K. R. Brown, 2012, Phys. Rev. A **86**, 042336.

Schoelkopf, R. J., and S. M. Girvin, 2008, Nature (London) **451**, 664.

Shinoda, M., 2002, J. Appl. Probab. **39**, 1.

Shor, P. W., 1996, in *the Proceedings of the 37th Annual Symposium on Foundations of Computer Science* (IEEE, Washington, DC) [http://dl.acm.org/citation.cfm?id=875509&CFID=684014415&CFTOKEN=51272421].

Shor, P. W., 1997, SIAM J. Comput. **26**, 1484.

Siva, K., and B. Yoshida, 2016, arXiv:1603.07805.

Spohn, H., 1977, Lett. Math. Phys. **2**, 33.

Stark, C., lode Pollet, A. Imamoğlu, and R. Renner, 2011, Phys. Rev. Lett. **107**, 030504.

Steane, A. M., 1996, Phys. Rev. Lett. **77**, 793.

Suchara, M., S. Bravyi, and B. Terhal, 2011, J. Phys. A **44**, 155301.

Temme, K., 2014, arXiv:1412.2858.

Temme, K., and M. J. Kastoryano, 2015, arXiv:1505.07811.

Terhal, B. M., 2015, Rev. Mod. Phys. **87**, 307.

Terhal, B. M., and G. Burkard, 2005, Phys. Rev. A **71**, 012336.

Trebst, S., P. Werner, M. Troyer, K. Shtengel, and C. Nayak, 2007, Phys. Rev. Lett. **98**, 070602.

Tsomokos, D. I., T. J. Osborne, and C. Castelnovo, 2011, Phys. Rev. B **83**, 075124.

Tupitsyn, I. S., A. Kitaev, N. V. Prokof'ev, and P. C. E. Stamp, 2010, Phys. Rev. B **82**, 085114.

Vezzani, A., 2003, J. Phys. A **36**, 1593.

Vidal, J., S. Dusuel, and K. P. Schmidt, 2009, Phys. Rev. B **79**, 033109.

Viyuela, O., A. Rivas, and M. A. Martin-Delagado, 2012, New J. Phys. **14**, 033044.

Walker, K., and Z. Wang, 2012, Front. Phys. **7**, 150.

Wang, C., J. Harrington, and J. Preskill, 2003, Ann. Phys. (Amsterdam) **303**, 31.

Watson, F. H. E., E. T. Campbell, H. Anwar, and D. E. Browne, 2015, arXiv:1503.08800.

Wecker, D., B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, 2014, Phys. Rev. A **90**, 022305.

Wegner, F. J., 1971, J. Math. Phys. (N.Y.) **12**, 2259.

Weiss, U., 2012, *Quantum Dissipative Systems* (World Scientific Publishing Co. Pte. Ltd., Singapore), Vol. 2.

Wen, X.-G., 2003, Phys. Rev. Lett. **90**, 016803.

Wen, X.-G., 2004, *Quantum Field Theory of Many-Body Systems* (Oxford Graduate Texts, New York).

Wilczek, F., 1982, Phys. Rev. Lett. **49**, 957.

Witten, E., 1988, Commun. Math. Phys. **117**, 353.

Wootton, J. R., 2012, J. Mod. Opt. **59**, 1717.

Wootton, J. R., 2013, Phys. Rev. A **88**, 062312.

Wootton, J. R., J. Burri, S. Iblisdir, and D. Loss, 2014, Phys. Rev. X **4**, 011051.

Wootton, J. R., V. Lahtinen, and J. K. Pachos, 2009, in *Proceedings from the 4th workshop on the Theory of Quantum Computation, Communication and Cryptography*, Lecture Notes in Computer Science, edited by Andrew Childs and Michele Mosca, Vol. 5906, p. 56 [http://link.springer.com/book/10.1007%2F978-3-642-10698-9].

Wootton, J. R., and D. Loss, 2012, Phys. Rev. Lett. **109**, 160503.

Wootton, J. R., and J. K. Pachos, 2011a, Phys. Rev. Lett. **107**, 030503.

Wootton, J. R., V. Lahtinen, B. Doucot, and J. K. Pachos, 2011b, Ann. Phys. (Amsterdam) **326**, 2307.

Yeomans, J. M., 1992, *Statistical Mechanics of Phase Transitions* (Oxford University Press, New York).

Yoshida, B., 2011, Ann. Phys. (Amsterdam) **326**, 2566.

Yoshida, B., 2013, Phys. Rev. B **88**, 125122.

Yoshida, B., 2014, arXiv:1404.0457.

Yoshida, B., and I. L. Chuang, 2010, Phys. Rev. A **81**, 052302.

Zanardi, P., 2000, Phys. Rev. A **63**, 012301.

Zanardi, P., and M. Rasetti, 1997, Phys. Rev. Lett. **79**, 3306.

Zeng, B., A. Cross, and I. L. Chuang, 2011, IEEE Trans. Inf. Theory **57**, 6272.

Zilberberg, O., B. Braunecker, and D. Loss, 2008, Phys. Rev. A **77**, 012327.