

Error Suppression and Error Correction in Adiabatic Quantum Computation: Techniques and Challenges

Kevin C. Young* and Mohan Sarovar

Scalable and Secure Systems Research (08961), Sandia National Laboratories, Livermore, California 94550, USA

Robin Blume-Kohout

Advanced Device Technologies (01425), Sandia National Laboratories, Albuquerque, New Mexico 87185, USA

(Received 10 September 2012; revised manuscript received 23 July 2013; published 13 November 2013)

Adiabatic quantum computation (AQC) has been lauded for its inherent robustness to control imperfections and relaxation effects. A considerable body of previous work, however, has shown AQC to be acutely sensitive to noise that causes excitations from the adiabatically evolving ground state. In this paper, we develop techniques to mitigate such noise, and then we point out and analyze some obstacles to further progress. First, we examine two known techniques that leverage quantum error-detecting codes to suppress noise and show that they are intimately related and may be analyzed within the same formalism. Next, we analyze the effectiveness of such error-suppression techniques in AQC, identify critical constraints on their performance, and conclude that large-scale, fault-tolerant AQC will require error correction, not merely suppression. Finally, we study the consequences of encoding AQC in quantum stabilizer codes and discover that generic AQC problem Hamiltonians rapidly convert physical errors into uncorrectable logical errors. We present several techniques to remedy this problem, but all of them require unphysical resources, suggesting that the adiabatic model of quantum computation may be fundamentally incompatible with stabilizer quantum error correction.

DOI: [10.1103/PhysRevX.3.041013](https://doi.org/10.1103/PhysRevX.3.041013)

Subject Areas: Quantum Physics, Quantum Information

I. INTRODUCTION

Adiabatic quantum computation (AQC) is expected to be inherently robust against certain errors, such as dephasing and energy relaxation [1,2]. This robustness suggests the possibility of an easier route to scalable quantum computation than the conventional gate-based “circuit” model, with less stringent requirements for fault tolerance and fewer resources devoted to error suppression and correction. However, AQC’s inherent robustness is not sufficient for fault tolerance. For example, several studies [2–6] show that single-qubit noise can drive undesirable transitions out of the adiabatic ground state. In response, error-suppression techniques have been developed that can reduce the rate at which these transitions occur. However, it is well understood that error suppression alone is insufficient for fault tolerance in the circuit model. Fault tolerance requires an additional mechanism to remove the entropy generated by errors that do occur in the encoded system—i.e., error correction. Since this thermodynamic argument is independent of the computational model, we reasonably expect that achieving fault-tolerant AQC will also require some form of error correction.

In this paper, we address both error suppression and error correction in AQC, and prove several facts about them. We begin our discussion in Sec. II with an overview of the failure mechanisms present in AQC. In Sec. III, we discuss error-suppression techniques based on error-detecting quantum stabilizer codes. Currently known suppression strategies include energy-gap protection (EGP) [7], in which the addition of the stabilizer generators to the system Hamiltonian causes errors to incur large energetic penalties; dynamical decoupling (DD)[8], whereby stabilizer generators are applied periodically as unitary operators, refocusing errors much like traditional spin echos; and Zeno effect suppression [9], which prevents errors from accumulating through frequent measurements of the stabilizer generators. These three techniques apparently operate by very different physical mechanisms. However, Facchi *et al.* [10] have shown that both Zeno suppression and DD may be viewed as limiting cases of a more general mathematical framework. We extend this unification by showing that DD and EGP may both be understood using the same formalism, and that the two methods are effectively equivalent in their error-suppression power. Our result shows that all three error-suppression techniques may be considered functionally equivalent for AQC, with some important caveats related to their physical implementation and their behavior in the presence of a thermal bath.

With Sec. IV, we turn our attention to error correction, where we discover two fundamental obstacles to the adaptation of stabilizer-code quantum error correction to

*kyoung@sandia.gov

Published by the American Physical Society under the terms of the [Creative Commons Attribution 3.0 License](https://creativecommons.org/licenses/by/3.0/). Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

AQC: (i) In the presence of the adiabatic Hamiltonian, correctable errors rapidly become uncorrectable errors, and (ii) the resources required by circuit-model error correction, such as stabilizer measurements and unitary gates, are generally outside the scope of the adiabatic paradigm. In Sec. V of the paper, we explore possible ways to overcome these obstacles when adapting stabilizer-based error correction to AQC. We identify alternative implementations of logical operators and continuous-time error correction by cooling as possible yet very challenging routes forward. Section VI finishes with a discussion of our view of the future of adiabatic fault tolerance.

The major results discussed in this manuscript are the following.

- (1) In the Hamiltonian formalism, the energy-gap protection and dynamical-decoupling methods for error suppression based on stabilizer encodings may be described by a unified mathematical formalism.
- (2) AQC appears to be fundamentally incompatible with stabilizer quantum error correction. Patching this incompatibility requires unphysical resources.

A companion paper entitled *Error Suppression and Correction in Adiabatic Quantum Computation: Non-equilibrium Dynamics* [11] develops a dynamical model for describing error suppression and correction in AQC, and discusses most of the results in this paper from a dynamical perspective.

II. THE QUANTUM ADIABATIC ALGORITHM AND IMPORTANT FAILURE MODES

The quantum adiabatic algorithm [1] operates by slowly changing an N -qubit system's Hamiltonian from a simple separable Hamiltonian H_{init} , whose ground state is easily prepared, to a final target Hamiltonian H_{prob} , whose highly nontrivial ground state encodes the solution to a problem of interest. The adiabatic theorem promises that if the Hamiltonian is changed slowly enough, then the system will remain in its (time-varying) ground state, and thus the solution can be read out by measuring the final state. In the simplest case, the closed-system dynamics may be described by a time-dependent Hamiltonian

$$H_{\text{AQC}}(t) = [1 - s(t)]H_{\text{init}} + s(t)H_{\text{prob}}, \quad (1)$$

where $s(0) = 0$, $s(T) = 1$, and T is the adiabatic interpolation time. More complicated interpolation schemes, including the addition of ancillary Hamiltonians, have been considered elsewhere [12,13] but are unnecessary for our discussions here. Importantly, the adiabatic model does not require the ability to perform high-quality quantum gates or measurements during the computation, which are key elements of fault-tolerant circuit-model quantum computation.

The adiabatic model is expected to be robust to some errors that plague other computational models, such as the cluster-state or circuit models. In particular, Ref. [2] showed that AQC possesses an inherent robustness to

both control errors and some forms of decoherence. In this section, we give a brief overview of what can go wrong with adiabatic computations and discuss which failure modes can be suppressed with current techniques.

Perhaps the best-known failure mode of AQC is diabatic errors (also known as Landau-Zener transitions), in which the Hamiltonian is varied too quickly and the state fails to track the instantaneous ground state. The obvious solution to this problem is to perform the interpolation more slowly, although identifying diabatic errors and determining the maximum allowable speed are nontrivial. For certain problems, the location of the minimum energy gap between the ground and first excited states is well known (e.g., for an adiabatic Grover search, it happens exactly in the middle [12]). This knowledge permits efficiently varying the interpolation speed so that it proceeds rapidly in regions where the gap is large and slowly in regions where the gap is small. Unfortunately, the location and magnitude of the minimum gap are not known for most problems. The blunt approach of slowing the entire interpolation is problematic because even simple problems may have exponentially small minimum gaps, leading to exponentially long interpolation times. This problem is not the focus of this paper, but solving it will necessarily involve (or enable) great leaps in our understanding of the computational complexity of the adiabatic algorithm.

While the system can be quite susceptible to diabatic transitions, AQC is known to possess some intrinsic robustness to Hamiltonian control errors [2]. As long as the evolution remains adiabatic, small perturbations to the intermediate Hamiltonians are likely to be unimportant. However, errors in the final Hamiltonian can be fatal, since it is the final Hamiltonian that encodes the problem to be solved. In particular, if the final Hamiltonian is close to a critical point, small perturbations may drastically alter the character of the ground state. This failure mode is addressed in Ref. [14].

Finally, because the success of an adiabatic interpolation relies ultimately on the population in the final ground state, it is robust to environmental couplings that cause decoherence in the eigenbasis of $H(s)$ (sometimes called dephasing). However, most system-bath couplings will cause transitions from the adiabatic ground state as well as Lamb shifts of the system Hamiltonian. The resulting open-system dynamics may be radically different from those of the ideal closed system. This paper is largely concerned with these errors, and we focus on techniques designed to suppress and correct the influence of these system-bath couplings in a manner consistent with the adiabatic paradigm. Specifically, we attempt to avoid reliance on quantum gates and measurements as much as possible.

III. SUPPRESSING ERRORS IN AQC

In this section, we introduce quantum stabilizer codes and discuss their role in protecting adiabatic quantum computations. A system undergoing adiabatic quantum

evolution while coupled to an external environment or bath is described in a tensor-product Hilbert space $\mathcal{H}_{\text{sys}} \otimes \mathcal{H}_{\text{env}}$ by a Hamiltonian

$$H(t) = H_{\text{AQC}}(t) + \sum_{j=1}^{n_e} E_j \otimes B_j + H_B. \quad (2)$$

Here, H_{AQC} acts on the system and performs the adiabatic evolution, and H_B is the bath Hamiltonian. B_j is a bath operator, and E_j is a single-qubit Pauli error operator. The total number of system-bath-coupling operators is n_e and is generally proportional to the total number of qubits in the system. Eigenstates of $H_{\text{AQC}}(t)$ may be labeled as $|n, k\rangle_t$ according to their principal quantum number n and an index k distinguishing any degeneracy; the subscript labeling the time is necessary because the Hamiltonian is time dependent (so $|n, k\rangle_t \neq |n, k\rangle_{t'}$). The system-bath-interaction terms in the Hamiltonian can cause the computation to fail by inducing transitions out of the adiabatically evolving ground state.

A. Quantum stabilizer codes

All currently known techniques for suppressing the errors induced by the system-bath-interaction terms rely on encoding the system in an error-detecting stabilizer code [15,16]. ‘‘Encoding’’ comprises

- (1) introducing (many) extra physical qubits to the system, and
- (2) mapping the original computational qubits (on which the computation is performed) into logical qubits that are distributed across many physical qubits, much as in a classical repetition code.

All the physical qubits together define a large system with a Hilbert space \mathcal{H}_{sys} . The logical qubits are a subsystem, corresponding to a factor space \mathcal{L} , so the entire Hilbert space factors as $\mathcal{H}_{\text{sys}} = \mathcal{L} \otimes \mathcal{S}$. The complementary subsystem \mathcal{S} is the syndrome subsystem. This factorization into subsystems is carefully chosen so that physical errors (on a small number of physical qubits) can be detected by Pauli measurements on the syndrome. These measurement operators, the stabilizer generators, are the quantum analogue of parity checks in classical linear block codes. The stabilizer generators generate an Abelian subgroup of the Pauli group (the stabilizer group, comprising all possible products of generators) and are used to compactly define the code. Stabilizers can be measured without disturbing the encoded quantum information because logical qubit operators (by definition) commute with the stabilizers.

An encoded system is always initialized in a known eigenstate of all the stabilizers. Subsequently, measuring the stabilizers will reveal (detect) any error operation that anticommutes with one or more stabilizer generators (since such a detectable error necessarily flips the sign of some stabilizer eigenvalues). These errors will be generated by terms in the system-bath Hamiltonian, whose effect on the

system is to apply E_j operations [Eq. (2)] on the system. Thus, we want to choose a code such that each E_j in the system-bath interaction anticommutes with at least one of the stabilizer generators. Utilizing a code with N_g stabilizer generators adds N_g physical qubits, enlarging the system’s Hilbert space by a factor of 2^{N_g} . The encoding process replaces the original problem Hamiltonian H_{AQC} with an encoded Hamiltonian, in which the Pauli operators σ_x , σ_y , and σ_z (that acted on physical bits in the unencoded Hamiltonian) are replaced by the code’s logical operators \bar{X} , \bar{Y} , and \bar{Z} . In addition, a time-dependent system control Hamiltonian H_C , expressible in terms of the code’s stabilizer generators, is added to implement any desired error suppression. The encoded Hamiltonian is then

$$\bar{H}(t) = \bar{H}_{\text{AQC}}(t) + H_C(t) + \sum_{j=1}^{N_g} E_j \otimes B_j + H_B. \quad (3)$$

We have assumed that the system-bath interaction remains qualitatively the same after the encoding, but is extended to $N_e > n_e$ terms to describe the extra qubits. Importantly, we have assumed that any controls we apply act only on the system and have no effect on system-bath couplings.

States of the encoded system may now be labeled by the same two quantum numbers as before, but with N_g additional binary quantum numbers, collected into the vector $\boldsymbol{\nu}$, indicating the eigenvalues of the stabilizer generators $S_j|n, k; \boldsymbol{\nu}\rangle_t = (-1)^{\nu_j}|n, k; \boldsymbol{\nu}\rangle_t$, where $S_j \in \mathcal{S}$ is a generator of the stabilizer group and $\nu_j = \{0, 1\}$ is the j th element of the vector $\boldsymbol{\nu}$. We shall refer to the subspace on which all the stabilizer eigenvalues are +1 (and therefore all $\nu_j = 0$) as the code space. States in the code space will therefore be labeled as $|n, k; \mathbf{0}\rangle_t$. The projector onto the code space is time invariant and expressible in terms of the stabilizer generators as

$$\mathbf{P}_0 = \frac{1}{2^{N_g}} \prod_{k=1}^{N_g} (1 + S_k). \quad (4)$$

Encoding will help to protect logical information (stored initially in the code space) by permitting active suppression of errors that cause transitions out of the code space. Errors that mix states within the code space are necessarily high weight [17] (and therefore, hopefully, unlikely). Encoding also makes it possible, in principle, to correct errors, by using the results of stabilizer measurements to detect and identify errors, then inverting them. However, such correction operations traditionally require resources (measurements and gates) that we have abjured.

In the following discussion, we will make extensive use of the toggling frame, a rotating frame of reference defined in terms of the control and bath Hamiltonians. Transformations to and from this frame are effected by the unitary operator

$$\mathcal{U}_C(t_1, t_2) = \exp\left(-i \int_{t_1}^{t_2} [H_C(s) + H_B] ds\right). \quad (5)$$

There is no need to time order the integral because $H_C(s)$ may be written entirely in terms of the stabilizer generators of the code, all of which are mutually commuting, and H_B acts on a different part of the Hilbert space. The following notation is used for an operator A in the toggling frame: $\tilde{A}(t) \equiv \mathcal{U}_C^\dagger(0, t)A\mathcal{U}_C(0, t)$, while states in the toggling frame are related to states in the Schrödinger picture by $|\tilde{\psi}\rangle_t = \mathcal{U}_C^\dagger(0, t)|\psi\rangle_t$. Evolution of states in this frame is generated by the toggling-frame Hamiltonian $\tilde{H}(t) \equiv \mathcal{U}_C^\dagger(0, t)[\tilde{H}(t) - H_C(t) - H_B]\mathcal{U}_C(0, t)$.

B. Dynamical decoupling

Dynamical decoupling is a well-known quantum control technique for suppressing errors produced by spurious terms in a system's Hamiltonian [18]. The methods were first applied to AQC in Ref. [8], with higher-order strategies shown to be particularly effective in Ref. [19]. In DD, the stabilizer generators are applied as unitary operations by manipulating the control Hamiltonian H_C . The sequence in which they are applied is given by the vector \mathbf{n} , at times given by $K(t) \in \mathbb{Z}$, so that at time t , the last operator applied to the system is $S_{\mathbf{n}_{K(t)}}$. The unitary operator defining the toggling frame may then be written as $\mathcal{U}_C^{\text{DD}}(t) = \prod_{j=0}^{K(t)} S_{\mathbf{n}_j}$. The operator $\mathcal{U}_C(t)$ is an element of the full stabilizer group and therefore commutes with $\tilde{H}_{\text{AQC}}(t)$. In the toggling frame, the Hamiltonian takes the form

$$\tilde{H}_{\text{DD}}(t) = \tilde{H}_{\text{AQC}}(t) + \sum_{j=1}^{N_e} \tilde{E}_j^{\text{DD}}(t) \otimes \tilde{B}_j(t).$$

Since E_j is a Pauli operator, it either commutes or anticommutes with each member of the stabilizer group, and we may write

$$\tilde{E}_j^{\text{DD}}(t) = \mathcal{U}_C^{\text{DD}\dagger}(t)E_j\mathcal{U}_C^{\text{DD}}(t) = (-1)^{p_j(t)}E_j, \quad (6)$$

where $p_j(t) = 0$ if $[E_j, \mathcal{U}_C^{\text{DD}}(t)] = 0$ and $p_j(t) = 1$ if $\{E_j, \mathcal{U}_C^{\text{DD}}(t)\} = 0$. A well-chosen DD sequence will cause $p_j(t)$ to rapidly alternate between 0 and 1, which modulates the system-environment coupling (in the toggling frame) by a rapidly oscillating function of t . The unitary operator governing the evolution at time t is $U(t, 0) = \exp_+ \{-i \int_0^t ds \tilde{H}_{\text{DD}}(s)\}$, and if the DD sequence is well chosen, the system-environment coupling averages to zero on time scales longer than the DD interpulse period [i.e., the integral vanishes thanks to the modulation factor $(-1)^{p_j(s)}$ in the exponential].

C. Energy-gap protection

The EGP approach, introduced in Ref. [7], appears quite different. It uses a constant-in-time control Hamiltonian, given by a sum of the stabilizer generators $H_C^{\text{EGP}}(t) = -\alpha \sum_{m=1}^{N_g} S_m$, with $\alpha > 0$. States in the code space are

then eigenstates of H_C with eigenvalue $-\alpha N_g$, but any state outside the code space is subjected to an energy penalty. Since H_C is a function only of the stabilizer generators, $\mathcal{U}_C^{\text{EGP}}(t)$ again commutes with the code's logical operators that comprise $\tilde{H}_{\text{AQC}}(t)$, so we can write the Hamiltonian in the toggling frame as

$$\tilde{H}_{\text{EGP}}(t) = \tilde{H}_{\text{AQC}}(t) + \sum_{j=1}^{N_e} \tilde{E}_j^{\text{EGP}}(t) \otimes \tilde{B}_j(t).$$

Error operators in the EGP toggling frame can be shown to take the form

$$\tilde{E}_j^{\text{EGP}}(t) = E_j e^{2i\alpha t \sum_{\{S_m, E_j\}=0} S_m} = e^{-2i\alpha t \sum_{\{S_m, E_j\}=0} S_m} E_j, \quad (7)$$

where the sums are taken over all stabilizer generators S_m that anticommute with the error operator E_j . To obtain this expression, we have exploited the following: (i) The stabilizer generators commute with each other (allowing easy manipulation of the exponential operators), and (ii) each generator either commutes or anticommutes with the noise operators: $S_m E_j = \pm E_j S_m$. Let w_j be the number of generators that anticommute with E_j . Then, the action of this toggling-frame Hamiltonian on any state $|\tilde{\Psi}\rangle = |\tilde{\psi}_c\rangle \otimes |\tilde{\phi}\rangle \in \mathcal{H}_{\text{sys}} \otimes \mathcal{H}_{\text{env}}$ with $|\tilde{\psi}_c\rangle$ in the code space is

$$\tilde{H}_{\text{EGP}}(t)|\tilde{\Psi}\rangle = \left(\tilde{H}_{\text{AQC}}(t) + \sum_{j=1}^{N_e} E_j e^{2iw_j\alpha t} \otimes \tilde{B}_j(t) \right) |\tilde{\Psi}\rangle. \quad (8)$$

Thus, the coupling term $E_j \otimes \tilde{B}_j(t)$ is modulated by a factor of $e^{2iw_j\alpha t}$. Just as in the case of DD (above), the error terms are modulated by an oscillating function in the interaction picture—which ensures that they average to zero on sufficiently long time scales as long as the frequency of oscillation is larger than the typical frequencies in $\tilde{B}_j(t)$. In the case of EGP, the oscillations are smooth and sinusoidal, whereas for (impulsive) DD, the oscillations are square waves in time. EGP can be made to mimic a decoupling sequence by choosing α so that the EGP oscillations match the frequency of a DD sequence. Numerical studies provide evidence that in such cases, EGP and DD suppress errors equally well, as shown in Fig. 1.

In fact, there is no requirement that the weights α be constant in time or equal across the stabilizer generators. Many dynamical-decoupling schemes vary the time interval between the pulses. For example, Uhrig's dynamical-decoupling scheme (UDD) [20] chooses the pulse arrival times as $t_n = T \cos[n\pi/2(N+1)]$, where N is the total number of pulses in the time interval τ . To mimic this UDD sequence, where the modulation frequency is not constant in time, we choose a time-dependent weight term $\alpha(t) = NT/\sqrt{t(T-t)}$. This approach was used, suitably regularized and in the context of a single qubit, by the authors of Ref. [21] to produce an effective UDD sequence using




FIG. 1. Logarithmic plot of the failure probabilities as a function of the total evolution time T for a two-logical-qubit AQC defined by the Hamiltonian $H(t) = (1 - t/T)(\sum_i X_i) + (t/T) \times (Z_1 + Z_2 + Z_1 Z_2)$ under the influence of Hamiltonian noise $H_\eta(t) = \sum_i \eta_i(t) Z_i$. In these numerical simulations, η_i is a classical stochastic process of $1/f$ type with spectrum $S(\omega) = 10^{-3}/\omega$, and the success probabilities are averaged over 1000 instances of this noise process for each evolution time T . Four simulations are performed: (i) the unencoded system (thick black line), (ii) the system encoded into four physical qubits using the $[[4, 2, 2]]$ quantum code [15] (dashed red line), (iii) the encoded system with EGP applied at strength $\Omega = 1$ (thin blue line), and (iv) the encoded system with DD applied at frequency $\Omega/2\pi$ (black crosses). The initial performance increase is due to reduction of diabatic errors, while later performance degradation is due to the accumulation of errors. Both EGP and DD are capable of suppressing this type of noise equally well, as shown by the approximately equal success probabilities of cases (iii) and (iv). Encoding without error suppression [case (ii)] performs especially poorly because twice as many qubits are exposed to noise as in the unencoded case and no measures are taken to suppress it.

continuous controls. More generally, allowing α to vary in time allows the strongest identification between the DD and EGP approaches and a unified treatment of both as quantum control protocols. For instance, choosing $\alpha_j(t) = \sum_i \pi \delta(t - t_i^j)/2$ applies S_j at time t_i^j as a unitary operator (impulsive DD), but in the EGP formalism, and points to how one might smoothly interpolate between the two methods. In this approach, optimal control techniques can be used to choose the $\alpha_j(t)$ to optimally mitigate the system-bath interaction.

D. A few important differences

The discussion above indicates that DD and EGP suppression methods have very similar behavior with respect to noise and are capable of providing approximately equivalent error suppression. This relationship is further examined in the context of filter functions in Appendix B and in a dynamical framework in Ref. [11]. But, while

these error-suppression techniques are closely related, it is important to be aware of some key differences, including the relative ease of their physical implementation and their effect on thermalization.

Many codes possess high-weight stabilizer generators. These codes cannot reasonably be implemented by EGP, since high-weight Hamiltonians are experimentally infeasible. They can be implemented by DD, however, by exploiting the fact that unitary operators may be generated by many different Hamiltonians. A high-weight DD pulse can be generated by single-body Hamiltonians, e.g.,

$$XXX = e^{-i(\pi/2)XXX} = e^{-i(\pi/2)(XII+IXI+IIX)}. \quad (9)$$

However, the single-body Hamiltonian implementing the DD does not commute with the encoded AQC Hamiltonian. If its implementation is not impulsive (i.e., the DD Hamiltonian is not significantly stronger than the AQC Hamiltonian), then it will be imperfect. The error can be computed fairly easily with the Baker-Campbell-Hausdorff formula and vanishes in the impulsive limit. Despite this complication, DD appears to be the only option for implementing high-weight stabilizer generators. In fact, there exist codes [22] that possess a large number of two-body stabilizer generators and only a few high-weight generators. These codes may benefit from a hybrid approach where low-weight generators are added as energy penalties, while those of high weight are included as DD pulses.

Another important difference between DD and EGP is their behavior in thermal environments. Whereas EGP establishes a real energy difference between the code space and the various syndrome spaces, DD does not. If a system with an EGP Hamiltonian is coupled to a cold thermal reservoir, thermalization will lead to a Gibbs distribution ($\rho \propto e^{-\beta H}$), in which the syndrome spaces are thermally populated according to their energy, and the code space is preferentially populated (since it has the lowest energy). In contrast, DD creates no real energy difference between syndrome spaces, and the system's steady-state population will be uniformly distributed across all the syndrome spaces. This distinction between EGP and DD will be important in Sec. V.

IV. ERROR CORRECTION

The error-suppression mechanisms described above reduce the rate at which errors appear in the system by effectively renormalizing the system-bath coupling. (See Ref. [11] for details.) But, the coupling cannot be eliminated completely without using an infinite amount of energy (e.g., DD pulses applied at infinitely high frequency, or an infinitely strong EGP Hamiltonian). Since our resources are finite, physical errors will still accumulate over long time scales and eventually cause logical errors. Typical AQC problems (e.g., combinatoric optimization) may have some intrinsic robustness to logical errors (e.g., a

few bits flipped in the solution to a Boolean satisfiability problem can be fixed efficiently with classical postprocessing), but as they accumulate, they will cause the computation to fail. Using large systems for long computations thus still requires some form of error correction.

Implementing error correction requires, to start, utilizing a code that is error correcting (rather than just error detecting). Not only must each error E_k anticommute with at least one stabilizer generator, but the set of generators with which it anticommutes (E_k 's syndrome) must uniquely identify E_k 's effect on the encoded computation, and therefore how to correct it. For simplicity, we will consider nondegenerate codes [23], in which each correctable error's syndrome is unique [24]. Such codes can thus determine what error has occurred (instead of simply detecting that some error occurred). Each correctable error can be labeled by its syndrome, a binary vector ν indicating which stabilizer generators anticommute with the error:

$$S_j E_\nu S_j E_\nu = (-1)^{\nu_j}. \quad (10)$$

For example, in a code with four stabilizer generators, an error labeled E_{0101} commutes with S_1 and S_3 but anticommutes with S_2 and S_4 . A code for which the number of correctable errors is exactly equal to the total number of nontrivial syndromes (i.e., $2^{N_s} - 1$) is called perfect [15,16].

A. A challenge for error correction

The simplest way to incorporate error correction into an AQC is to just do the encoding, and nothing else. Perform the encoded adiabatic interpolation, then measure both the stabilizers and the logical Z operators on the final state. In the absence of errors, the stabilizer measurements would all yield +1, and the logical Z measurements would yield the answer to the computation. If one or more physical errors do occur and accumulate to produce a correctable error, then the correct answer can be decoded (classically) from the final measurement outcomes. So, this naïve implementation of error correction offers at least one advantage over pure error suppression.

However, it is unlikely to be sufficient. In the absence of some correcting mechanism during the evolution, errors will accumulate. With a probability that approaches 1 as the computation grows (in size and time), a logical error will occur, and the (decoded) final answer will be wrong. Thus, we anticipate an additional need for some ongoing entropy-extracting process.

First, however, we must address a different (yet arguably more pernicious) problem, which arises even for very small systems where bath-induced uncorrectable errors are rare: The encoded problem Hamiltonian transforms correctable errors to logical errors.

To illustrate this issue, we consider a simple case. Suppose the system is encoded and initialized in the ground state $|0\rangle$ of an initial Hamiltonian and evolves unperturbed under the adiabatically changing

Hamiltonian until, at time τ , a correctable Pauli error E_ν occurs. Then, the system evolves unperturbed through the end of the AQC, at which point we measure the code stabilizers. Because the Hamiltonian always commutes with the stabilizers and only a single correctable error has occurred, one might think that the error can be detected and identified by its syndrome ν , enabling a restoration of the system to the ground state (in the code space) by an application of E_ν . Unfortunately, in the time span between the error and its subsequent correction, things go horribly awry.

The overall evolution according to our simplified error model is

$$|\psi\rangle_T = E_\nu \mathcal{U}_{\text{AQC}}(\tau, T) E_\nu \mathcal{U}_{\text{AQC}}(0, \tau) |0; \mathbf{0}\rangle_0, \quad (11)$$

where the unitary evolution generated by the adiabatic Hamiltonian is given by a time-ordered exponential

$$\mathcal{U}_{\text{AQC}}(\tau, T) = \exp_+ \left(-i \int_\tau^T \bar{H}_{\text{AQC}}(s) ds \right). \quad (12)$$

(We neglect the as-yet-unspecified error-suppressing control Hamiltonian, as its presence does not change the result.) We assume that the only error that occurs is E_ν and that there are not other deviations (such as Landau-Zener transitions) from ideal adiabatic evolution, so $\mathcal{U}_{\text{AQC}}(0, \tau) |0; \mathbf{0}\rangle_0 = |0; \mathbf{0}\rangle_\tau$. Now, the encoded AQC Hamiltonian $\bar{H}_{\text{AQC}}(s)$ is a weighted sum of the code's logical X , Y , and Z operators. Each is a Pauli operator, so it either commutes or anticommutes with the error operator E_ν . The encoded AQC Hamiltonian (at any normalized time s) splits into a commuting and an anti-commuting term

$$\bar{H}_{\text{AQC}}(s) = \bar{H}_\nu^+(s) + \bar{H}_\nu^-(s), \quad (13)$$

where $[\bar{H}_\nu^+(s), E_\nu] = \{\bar{H}_\nu^-(s), E_\nu\} = 0$. After some algebra, Eq. (11) becomes

$$|\psi\rangle_T = \exp_+ \left(-i \int_\tau^T [\bar{H}_\nu^+(s) - \bar{H}_\nu^-(s)] ds \right) |0; \mathbf{0}\rangle_\tau. \quad (14)$$

Between the time when the error happens (τ) and when it is corrected (T), the encoded system experiences a new, effective Hamiltonian $\bar{H}'_\nu = \bar{H}_\nu^+(s) - \bar{H}_\nu^-(s)$. Since the state $|0; \mathbf{0}\rangle_\tau$ is not generally an eigenstate of \bar{H}'_ν , the system will undergo unintended evolution within the code space, moving it out of the ground state (a logical error).

The root problem here is that the encoded adiabatic Hamiltonian acts differently on different syndrome spaces (eigenspaces of the stabilizers). Correctable errors flip various stabilizers' eigenvalues, moving the system from one syndrome space to another. Only on the code space itself is \bar{H}_{AQC} guaranteed to act like the original problem Hamiltonian H_{AQC} . In other words, the AQC Hamiltonian itself rapidly turns correctable errors into logical errors. Like nontransversal implementations of logic gates in the circuit model, naïvely encoded AQC Hamiltonians cause

errors to propagate, and decoding at the end of the computation is unlikely to be effective.

The problem persists even if errors are corrected during the computation. Even the best error correction is not instantaneous, so errors will survive for some time before being corrected. In large computations, the equilibrium between noise and correction occurs at a finite density of errors, meaning that at every instant, the system will be out of the code space. Thus, it is absolutely necessary to modify the encoded Hamiltonian so that \bar{H}_{AQC} is sufficiently similar to H_{AQC} , not just on the code space but on every likely syndrome space.

B. Protected Hamiltonians

The promotion of physical errors to logical errors by the adiabatic Hamiltonian poses a serious threat to error correction in AQC using stabilizer codes. In this section, however, we demonstrate that this threat may be avoided by carefully choosing the logical operators—but doing so comes at a steep cost: These protected logical operators are complicated sums of high-weight Pauli operators.

Physical errors become logical errors because \bar{H}_{AQC} does not act identically on all of the syndrome spaces. In principle, it would be enough if any error mapped the ground state(s) of \bar{H}_{AQC} to another eigenstate of \bar{H}_{AQC}

$$\bar{H}_{\text{AQC}}(t)E_{\nu}|0, k; \mathbf{0}\rangle_t = \epsilon_{0, k; \nu}(t)E_{\nu}|0, k; \mathbf{0}\rangle_t.$$

But, our error-correcting code is not allowed to “know” what the ground state is. This condition has to hold for any encoded logical Hamiltonian built as a sum of logical operators. So, whenever an AQC is encoded in a stabilizer code, some logical Hamiltonians will rapidly transform some physical errors into logical errors unless the following stronger condition is satisfied: For any error E_{ν} and any eigenstate $|n, k\rangle$ of the logical Hamiltonian, $E_{\nu}|n, k; \mathbf{0}\rangle$ must be an eigenstate of the encoded logical Hamiltonian. In other words,

$$\bar{H}_{\text{AQC}}(t)E_{\nu}|n, k; \mathbf{0}\rangle_t = \epsilon_{n, k; \nu}(t)E_{\nu}|n, k; \mathbf{0}\rangle_t. \quad (15)$$

If this condition holds, then at least adiabatic evolution is still possible in the presence of errors (although errors might reduce or eliminate the gap, or even cause an excited state to become the ground state).

If we want Eq. (15) to hold, then the encoded Hamiltonian (and therefore the encoded logical operators) is constrained. The eigenstates specified in Eq. (15) form a complete orthonormal basis, so the most general encoded Hamiltonian that satisfies Eq. (15) is a sum of their projectors

$$\begin{aligned} \bar{H}_{\text{AQC}}(t) &= \sum_{n, k, \nu} \epsilon_{n, k; \nu}(t) |n, k; \nu\rangle_t \langle n, k; \nu| \\ &= \sum_{n, k, \nu} \epsilon_{n, k; \nu}(t) E_{\nu} |n, k; \mathbf{0}\rangle_t \langle n, k; \mathbf{0}| E_{\nu}^{\dagger}. \end{aligned} \quad (16)$$

The projector $|n, k; \mathbf{0}\rangle_t \langle n, k; \mathbf{0}|$ may be written as a product of the projector onto the code space (expressible entirely in terms of the code’s stabilizers) and the logical projector onto the encoded state $|n, k\rangle_t$ (expressible entirely in terms of the code’s logical operators).

To construct an encoded Hamiltonian with these properties, we can exploit the freedom in defining logical operators of a code. Multiplying a logical operator by a stabilizer operator yields an equally valid representation of the same logical operator. We can represent elements of the stabilizer group by binary vectors \mathbf{j} , of length N_g , in which a 1 indicates the presence of the corresponding generator, i.e., $S_{\mathbf{j}} = S_1^{j_1} S_2^{j_2} \cdots S_{N_g}^{j_{N_g}}$. For example, if there are four stabilizer generators, then $S_{0110} \equiv S_2 S_3$. In this notation, the fact that (Pauli) errors either commute or anticommute with each of the stabilizers can be written as

$$S_{\mathbf{j}} E_{\nu} = (-1)^{j_{\nu}} E_{\nu} S_{\mathbf{j}}. \quad (17)$$

Then, if \bar{L}_i is a particular encoded logical operator, each stabilizer operator $S_{\mathbf{j}}$ defines an equivalent logical operator $\bar{L}_i S_{\mathbf{j}}$ that acts identically on the code space. Linear combinations, e.g., $\sum_{\mathbf{j}} \beta_{i\mathbf{j}} \bar{L}_i S_{\mathbf{j}}$, are also valid logical operators.

Now, any problem Hamiltonian can be written as a sum of logical operators

$$H_{\text{AQC}}(t) = \sum_i \alpha_i(t) L_i.$$

When we encode it, we can choose any encoding $L_i \rightarrow \sum_{\mathbf{j}} \beta_{i\mathbf{j}}(t) \bar{L}_i S_{\mathbf{j}}$ (where the \mathbf{j} sum is taken over all binary vectors of length N_g , and \bar{L}_i is an arbitrary encoded representation of L_i) as long as $\sum_{\mathbf{j}} \beta_{i\mathbf{j}}(t) = 1$. So, an equivalent encoded problem Hamiltonian is

$$\bar{H}_{\text{AQC}}(t) = \sum_i \alpha_i(t) \bar{L}_i \sum_{\mathbf{j}} \beta_{i\mathbf{j}}(t) S_{\mathbf{j}}. \quad (18)$$

Now, we impose the constraint that it must satisfy Eq. (15). Inserting Eq. (18) into Eq. (15) and multiplying by E_{ν} on the left, we obtain

$$\sum_i \alpha_i(t) \bar{L}_i \left(\sum_{\mathbf{j}} \beta_{i\mathbf{j}}(t) \Gamma_{i\mathbf{j}}^{\nu} \right) |n, k; \mathbf{0}\rangle_t = \epsilon_{n, k; \nu}(t) |n, k; \mathbf{0}\rangle_t,$$

where $E_{\nu} \bar{L}_i = \gamma_{i\nu} \bar{L}_i E_{\nu}$, and $\Gamma_{i\mathbf{j}}^{\nu} = \gamma_{i\nu} (-1)^{j_{\nu}}$ is 1 if E_{ν} commutes with $\bar{L}_i S_{\mathbf{j}}$ and -1 otherwise. Direct computation then shows that $\Gamma_{i\mathbf{j}}^{\nu}$ obeys the orthogonality condition

$$\sum_{\nu} \Gamma_{i\mathbf{j}}^{\nu} \Gamma_{ik}^{\nu} = 2^{N_g} \delta_{\mathbf{j}\mathbf{k}}, \quad \forall i. \quad (19)$$

Recall that $|n, k; \mathbf{0}\rangle_t$ is defined to be an eigenstate of $\sum_i \alpha_i(t) \bar{L}_i$ for any choice of $\alpha_i(t)$. Therefore, the sum $\sum_{\mathbf{j}} \beta_{i\mathbf{j}}(t) \Gamma_{i\mathbf{j}}^{\nu}$ cannot depend on i , and we may substitute

$$\lambda_{\nu}(t) = \sum_{\mathbf{j}} \beta_{i\mathbf{j}}(t) \Gamma_{i\mathbf{j}}^{\nu}. \quad (20)$$

With this definition, we see that $\epsilon_{n,k;\nu}(t) = \epsilon_{n,k;0}(t)\lambda_\nu(t)$, meaning that an encoded operator satisfying Eq. (15) acts identically (up to a scale factor λ_ν) on each syndrome subspace, including the code space. Correctable errors map eigenstates of \bar{H}_{AQC} in the code space to eigenstates in the appropriate syndrome space. In particular, the ground state maps either to the lowest-energy eigenstate in the syndrome space (if λ_ν is positive) or the highest-energy eigenstate (if λ_ν is negative). We refer to λ_ν as a deformation factor in Ref. [11], since it represents the deformation of the code-space energy structure when reproduced within syndrome spaces.

The most general Hamiltonian satisfying condition Eq. (15) can therefore be constructed by symmetrizing over errors E_ν :

$$\bar{H}_{\text{AQC}}^\lambda(t) = \sum_\nu \lambda_\nu(t) E_\nu \bar{H}(t) \mathbf{P}_0 E_\nu, \quad (21)$$

where $\bar{H}(t)$ is a Hamiltonian encoded using any logical operators. We refer to such Hamiltonians as protected Hamiltonians since adiabatic evolution under them does not lead to propagation of errors in the correctable syndrome subspaces. Furthermore, we define the canonical protected Hamiltonian by setting all λ_ν in Eq. (21) to 1

$$\bar{H}_{\text{AQC}}^p(t) = \sum_\nu E_\nu \bar{H}(t) \mathbf{P}_0 E_\nu. \quad (22)$$

The canonical protected Hamiltonian acts truly identically on every syndrome space, with no scale factor at all. Similar Hamiltonians were introduced in a slightly different context in Ref. [25].

Unfortunately, implementing the canonical protected logical operators seems infeasible. They generally contain (many) terms of very high weight, up to the total number of qubits in the system. However, if we allow a scale factor (i.e., choose noncanonical but still protected logical operators), then it is possible to reduce the maximum weight by choosing the coefficients λ_ν to eliminate the highest-weight logical operators from Eq. (21). Inverting Eq. (20), we have

$$\beta_{ij}(t) = \sum_\nu \frac{1}{2^{N_s}} \lambda_\nu(t) \Gamma_{ij}^\nu, \quad (23)$$

which can be used to choose the $\lambda_\nu(t)$'s such that the $\beta_{ij}(t)$'s corresponding to high-weight operators vanish. Note that we have explicitly included the time dependence in $\lambda_\nu(t)$ and $\beta_{ij}(t)$ to maintain generality, but it may be convenient to force them to be time independent, as in the canonical protected Hamiltonian [where $\lambda_\nu(t) = 1$]. The time dependence may be useful for constructing more sophisticated error-suppression schemes, such as increasing the penalty if a particular error is otherwise more likely at a certain time, but we do not consider such schemes here.

In Appendix A, we examine the [[5, 1, 3]] code [15] and show that while the canonical protected logical operators

are weight five, we can construct protected logical operators with weight three. Any logical operator for a distance- d code must have a weight of at least d , so it is encouraging that protected operators of weight d exist in this case. If such a construction is possible for any code (an open question), it would remove one obstacle to implementation.

However, a further challenge is the sheer number of Pauli operator terms required to implement these protected logical operators; for example, the canonical protected logical operators for n -qubit codes are sums of $O(2^n)$ Pauli operators. The minimum-weight protected logical operators that we construct in Appendix A are simpler but still contain $O(2^d)$ distinct Pauli operators. Applying a single weight- d operator (for $d \gg 1$) as a Hamiltonian is already challenging. (See the concluding discussion.) Balancing many such terms seems fantastically difficult. We suspect it will be feasible only if there exist protected logical operators in which the required sum of Pauli operators can be factored or otherwise decomposed into a sum of a few products. For example, $(X + Z)^{\otimes n}$, expanded as a sum of Pauli operators, contains 2^n terms—but because it factors, it is no harder to implement than $X^{\otimes n}$. We are not aware of any such structure in protected logical operators, but further research might reveal one.

V. ERROR CORRECTION BY LOCAL COOLING

Suppose that a way is found to prevent the adiabatic Hamiltonian from converting physical errors to logical errors (e.g., by implementing protected logical Hamiltonians). AQC would then face “only” the same problem that confronts circuit-model computation; errors accumulate over time. These errors are the manifestation of entropy injected into the system by coupling to the bath, and unless we actively pump that entropy out, the computation is likely to fail within a relatively short time. Active error correction, however, requires fast gates and high-weight stabilizer measurements during the computation—both of which are outside the standard AQC toolset (and infeasible in many of the physical systems on which AQC might be implemented). Local cooling offers a potential route around this difficulty.

In the local cooling model, each physical qubit is coupled to a very low-temperature bath that serves as the entropy sink for the system. The coupling is designed such that, if the stabilizer generators are added to the Hamiltonian, the bath is able to detect the increase in energy associated with an error and then to absorb that energy and reverse the error. The dynamics of error correction by cooling is worked out in detail in Ref. [11], and we sketch the scheme here.

We will consider the Hamiltonian of Eq. (3) with an additional local coupling of each error operator to a cold, damped reservoir given by

$$H_{S-R} = \sum_j E_j \otimes \sum_k g_k^{(j)} (b_k^{(j)} + b_k^{\dagger(j)}).$$

In addition to this interaction Hamiltonian, the reservoirs have free Hamiltonians of independent oscillators: $H_R = \sum_j \sum_k \omega_k^{(j)} b_k^{\dagger(j)} b_k^{(j)}$. We assume that each reservoir has a broad frequency distribution and is well damped, so that it is in its ground state at all times with high probability. It is shown in Ref. [11] that for such a system-reservoir coupling, a control Hamiltonian of the EGP type and a protected Hamiltonian implementation of the logical Hamiltonian are sufficient for an automated implementation of error correction where the excitations induced by local Pauli perturbations from the bath are quenched by the cold reservoir.

However, because the reservoir is coupled to the system through low-weight Pauli operators (E_j), it can only cool local excitations. Cooling away a high-weight error would have to be accomplished through a sequence of single-Pauli operations that reduce the error weight until the system is returned to the code space. However, as shown in Ref. [11], the cooling dynamics is a biased random walk in syndrome space, and therefore, in order for such a sequence of single-Pauli operations to successively cool away several errors, the energy landscape of the system must be such that the energy penalty associated with an error increases with its weight. That is, the energy of states in a correctable syndrome subspace must increase monotonically with the weight of the error that takes the code space to that syndrome space. The EGP control Hamiltonians corresponding to most stabilizer codes do not have this property. (We shall discuss exceptions in a moment.) For example, for the Abelian toric code [26], a state with two errors can have the same energy penalty as a state with one error if the two errors are neighboring; i.e., a string excitation does not have an excited-state energy proportional to the string length, and a local measurement of energy cannot distinguish between such degenerate errors. Such models have no “string tension” [27].

EGP control Hamiltonians that provide a favorable landscape, where syndrome-subspace energies scale with error weight, may be constructed explicitly as

$$H_C^{\text{EGP}} = \sum_k \sum_{\substack{E_p \text{ s.t.} \\ w(E_p)=k}} \delta(k) E_p \mathbf{P}_0 E_p, \quad (24)$$

where $\delta(k) > 0$ is the energy penalty associated with weight- k errors $\delta(k) < \delta(k + 1)$, and the second sum is over all E_p such that $w(E_p) = k$. However, as in the case of the protected Hamiltonian for logical evolution, such constructions result in Hamiltonians that are very high weight. In fact, recent work has shown that there are significant obstacles to constructing systems with the kind of energy landscape discussed above, known as a self-correcting quantum memory, using local stabilizers in two dimensions [28,29]. However, these results do not

rule out the construction of an energy landscape for which low-weight operators are penalized according to their weight but higher-weight operators are not. Such a code would exhibit string tension that “snaps” after the string length grows too long and would likely provide enhanced protection over standard codes. It might be constructed as in Eq. (24) by choosing the energy penalties so that high-weight stabilizers cancel, as we did in Sec. IV B to construct the protected logical operators. Furthermore, although the restriction to two-dimensional (planar) codes may be appealing from an engineering perspective, self-correcting lattice codes in four (and perhaps three) spatial dimensions do exist [30–34]. Embedding such higher-dimensional codes into two dimensions requires nonplanar connectivity, but this is not necessarily unrealistic; superconducting qubit systems routinely construct nonplanar interaction graphs by coupling distant qubits with wires [35].

VI. DISCUSSION

This work arose from our attempts to answer the question, “Can AQC be made fault tolerant?” We began without a clear picture of what “fault tolerance” should mean in the adiabatic context. Taking cues from the theory of circuit-model fault tolerance, we believed that developing a clear understanding of error suppression and error correction in AQC would be a necessary first step, whatever the ultimate definition turned out to be. In this paper, we have presented such a framework, investigating relationships between error-suppression methods and describing a serious challenge to adiabatic error correction. Through the methods of protected Hamiltonians and local cooling, we have even suggested techniques for avoiding these obstacles and correcting errors without resorting to circuit-model gates or syndrome measurements.

However, our analysis falls short of establishing a threshold theorem for adiabatic fault tolerance. We cannot prove fault tolerance because controlling the encoded AQC (using slowly varying Hamiltonians rather than fast gates) seems to require high-weight Hamiltonians. These Hamiltonians are not available in any feasible technology, to our knowledge. We are therefore unable to propose a feasible control protocol and so have no credible model to describe control errors. Without a plausible error model, we cannot attempt to prove fault tolerance or calculate a threshold. If in the future the implementation of protected logical Hamiltonians is shown to be practical, then it may become possible to construct fault-tolerant logical-operation protocols and to prove the existence of a threshold. But, absent such a breakthrough, logical Hamiltonians on encoded qubits appear to require unphysical resources, leaving us pessimistic that any form of fault tolerance will ever be achieved in a purely adiabatic model of quantum computation.

This conclusion rests on several assumptions. We now list these assumptions, and the limitations in our analysis, in the hope that future work may find ways to circumvent the obstacles identified here.

First, we have assumed that protecting AQC from noise will require the use of high-distance stabilizer codes. This assumption is motivated by fault tolerance in the circuit model, which depends critically on the use of high-distance quantum codes, because the computation must be robust to a constant (albeit low) density of errors. In the adiabatic setting, high-distance codes lead directly to two specific problems: (i) The encoded logical Hamiltonian quickly transforms correctable errors into uncorrectable logical errors, and (ii) the logical operators that comprise the encoded Hamiltonian necessarily have high weight. Our work suggests a solution to problem (i), but problem (ii) poses a greater challenge. In the circuit model, high-weight logical operators are implemented by performing many one- or two-qubit unitaries in parallel. Although such sequences are more complicated than unencoded gate operations, the (linear in N) increase in gate complexity is far outweighed by enhanced (exponential in N) resilience against noise. In the adiabatic model, however, logical operations are implemented as Hamiltonians—not unitaries. And, whereas high-weight unitaries can be implemented in $O(1)$ time using parallel gates, there is no comparable way to apply high-weight Hamiltonians. Solving this issue will likely require significant advances. Perturbative gadgets can approximate high-weight Hamiltonians with only weight-two interactions but introduce unprotected gauge qubits to the system. Furthermore, the published analyses [36] of gadget perturbation theory require couplings that scale exponentially with the operator weight, although perhaps more sophisticated gadget perturbation theories can be developed that reduce this penalty. Even more desirable (although correspondingly less likely) would be the development of qubits whose dominant interactions are naturally high weight [37]. Should this obstacle be overcome, our construction of protected logical operators (Sec. IV B) will be highly relevant.

Second, we have assumed that the only available error-correction mechanism is local cooling. The local nature of the cooling is clearly physically motivated, but local cooling can only drive single qubits and act upon local information [11]. It is possible that a more sophisticated cooling mechanism that acts on multiple qubits in a neighborhood (similar to a continuous-time error-decoding algorithm) could be constructed. Such a mechanism might obviate the need for a monotonic energy landscape, as stated in Sec. V. However, we have no constructive ideas for implementing such a cooling mechanism at this time.

Finally, an overriding assumption in this whole work is that elements of the circuit model, like fast gates and measurements, are not available. Of course, one could begin incorporating more elements of the circuit model

in order to implement error correction. However, in that case, the model of computation begins to look more and more like the circuit model itself. If fault-tolerant AQC demands development of all the resources required for computation in the circuit model, why bother with AQC at all?

ACKNOWLEDGMENTS

We acknowledge important discussions with Sandia’s AQUARIUS Architecture Team, especially those with Andrew Landahl and Anand Ganti. This work was supported by the Laboratory Directed Research and Development Program at Sandia National Laboratories. Sandia is a multiprogram laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the United States Department of Energy’s National Nuclear Security Administration under Contract No. DE-AC04-94AL85000.

APPENDIX A: SUBSYSTEM STRUCTURE OF PROTECTED HAMILTONIAN

The protected Hamiltonians defined in Eq. (21) naturally induce tensor-product decomposition on the system Hilbert space into a logical subsystem \mathcal{H}_{\log} and a syndrome subsystem $\mathcal{H}_{\text{synd}}$:

$$\bar{\mathcal{H}}_{\text{sys}} = (\mathcal{H}_{\log} \otimes \mathcal{H}_{\text{synd}}) \oplus \mathcal{H}_{\text{uncorr}} \quad (\text{A1})$$

Correctable eigenstates of the system $|n, k; \nu\rangle_t$ may be written in this context as $|n, k\rangle \otimes |\nu\rangle_t$. If the Hamiltonian is encoded as in Eq. (22), then its action on any correctable state is

$$\bar{H}_{\text{AQC}}^p(t)|n, k\rangle \otimes |\nu\rangle_t = \epsilon_{n,k}|n, k\rangle \otimes |\nu\rangle_t \quad (\text{A2})$$

Thus, the protected logical operators themselves inherit this tensor-product structure when acting on correctable states:

$$\bar{L}^p = \sum_{\nu} E_{\nu} \bar{L} P_0 E_{\nu} \rightarrow L \otimes \mathbf{I}_{n-k} \quad (\text{A3})$$

where \bar{L} is any logical operator of the code, and \mathbf{I}_{n-k} is the identity operator on the 2^{n-k} -dimensional syndrome space. Assuming a perfect code, a set of operators may be constructed that acts as Pauli operators on the syndrome bits. The stabilizer generators of the code may be interpreted as the Pauli Z operators on the syndrome bits:

$$S_j \rightarrow I \otimes Z_j \quad (\text{A4})$$

Pauli X operators flip stabilizer bits without introducing a phase and so may be constructed as

$$\sum_{\nu} E_{\nu \oplus \mathbf{j}} P_0 E_{\nu} \rightarrow I \otimes X_j \quad (\text{A5})$$

And, the Pauli Y operators can be constructed using the Pauli relation $iXZ = Y$:

$$i \sum_{\nu} E_{\nu} \mathbf{P}_0 E_{\nu} S_j \rightarrow I \otimes Y_i. \quad (\text{A6})$$

These operators now allow us to represent any physical Pauli operator in terms of its action on the logical and syndrome degrees of freedom. Any logical operator of the code L may be written as $L \rightarrow L \otimes O$ for some operator O that acts entirely on the syndrome space. Because the protected operators take the form $L \otimes I$, we may determine O by multiplication:

$$L^p L \rightarrow (L \otimes I)(L \otimes O) = I \otimes O. \quad (\text{A7})$$

Using this decomposition, one may easily see why sums of logical operators do not necessarily act on syndrome subspaces in the same way as they do on the code space: Their associated syndrome operator is different.

1. Example: $[[5, 1, 3]]$ code

Consider the quantum $[[5, 1, 3]]$ code, defined in terms of the stabilizer generators

$$S_1 = IXZZX, \quad (\text{A8})$$

$$S_2 = XIXZZ, \quad (\text{A9})$$

$$S_3 = ZXIXZ, \quad (\text{A10})$$

$$S_4 = ZZXIX \quad (\text{A11})$$

and logical operators

$$\bar{X} = XXXXX, \quad (\text{A12})$$

$$\bar{Z} = ZZZZZ. \quad (\text{A13})$$

In the tensor-product basis defined above, these logical operators take the form

$$\begin{aligned} \bar{X} \rightarrow X \otimes \frac{1}{4} &(-IIII + IIIZ + IIZI + IIZZ + IZII \\ &- IZIZ + IZZI + IZZZ + ZIII - ZIIZ - ZIZI \\ &- ZIZZ + ZZII - ZZIZ + ZZZI + ZZZZ), \end{aligned}$$

$$\begin{aligned} \bar{Z} \rightarrow Z \otimes \frac{1}{4} &(-IIII + IIIZ + IIZI - IIZZ + IZII + IZIZ \\ &- IZZI - IZZZ + ZIII + ZIIZ + ZIZI \\ &+ ZIZZ - ZZII + ZZIZ - ZZZI + ZZZZ). \end{aligned}$$

The syndrome parts of these operators are different and thus will behave differently on each syndrome subspace. To correct this behavior, we introduce the protected operators, defined by Eq. (22), which in this case take the form

$$\begin{aligned} \bar{X}^p &= -IIZXZ - IXIYY - IYYIX - IZXZI - XIYYI \\ &\quad - XXXXX - XYZZY - XZIIZ - YIXIY \\ &\quad - YXYZZ - YYIXI - YZZYX - ZIIZX \\ &\quad - ZXZII - ZYXYZ - ZZYXY, \\ \bar{Z}^p &= -IIZZY - IXXIZ - IYZYI - IZIXX - XIZIX \\ &\quad - XXIZI - XYYXZ - XZXYY - YIIYZ \\ &\quad - YXZXY - YYXZX - YZYII - ZIXXI \\ &\quad - ZXYYX - ZYIIY - ZZZZZ. \end{aligned}$$

These operators in the tensor-product basis are simply $X^p \rightarrow X \otimes I$ and $Z^p \rightarrow Z \otimes I$. These operators act consistently but require the use of high-weight operators. However, if we add up the minimum-weight versions of each logical operator, we have

$$\begin{aligned} \bar{X}_3 &= -IIZXZ - IXIYY - IYYIX - IZXZI - XIYYI \\ &\quad - XZIIZ - YIXIY - YYIXI - ZIIZX - ZXZII \\ &\rightarrow X \otimes (2\mathbf{P}_0 + \mathbf{I}_4/2), \\ \bar{Z}_3 &= -IIZXZ - IXIYY - IYYIX - IZXZI - XIYYI \\ &\quad - XZIIZ - YIXIY - YYIXI - ZIIZX - ZXZII \\ &\rightarrow Z \otimes (2\mathbf{P}_0 + \mathbf{I}_4/2). \end{aligned}$$

Because the syndrome part of these operators is the same, these operators act consistently across subspaces, avoiding the problems we described above. It remains an open question whether such low-weight constructions exist for higher-distance codes.

APPENDIX B. GENERALIZED FILTER FUNCTIONS FOR THE SINGLE-QUBIT CASE

The relationship between DD and EGP is nicely illustrated by a single-qubit example. We shall consider a single-qubit evolution in the presence of pure dephasing noise. Note that for a single qubit, the error-detection code allows us to stabilize a single state ($|+\rangle$ or $|-\rangle$ in the example below) rather than a subspace. The Hamiltonian describing noisy qubit evolution is

$$H(t) = \frac{1}{2}c(t)\sigma_x + \frac{1}{2}\eta(t)\sigma_z. \quad (\text{B1})$$

Here, $c(t)$ is the control field that could either implement DD or EGP, and $\eta(t)$ is the stochastic noise. Dynamical decoupling proceeds by applying a sequence of X -type pulses to perturbatively decouple the noise, while EGP introduces a time-independent energy penalty for noise-induced error transitions. We consider a classical approximation of the system-bath coupling and represent the bath fluctuations as a classical stochastic process. Taking this simplification is not strictly necessary for what follows, but we do so for clarity. In the toggling frame, this Hamiltonian takes the form




FIG. 2. The filter functions for periodic dynamical decoupling (solid line) and continuous amplitude driving (dashed line). The filter minima can be made to overlap at first order in the Magnus expansion.

$$\tilde{H}(t) = \frac{1}{2} \eta(t) \sigma_z \exp\left(-2i \int_0^t c(s) ds \sigma_x\right). \quad (\text{B2})$$

Defining the fidelity as the probability of finding the state in the $|+\rangle$ state after the evolution,

$$\mathcal{F}(\tau) = \langle \langle + | U(\tau) | + \rangle \rangle_\eta,$$

where $\langle \cdot \rangle_\eta$ indicates a classical stochastic average over instances of the noise $\eta(t)$. To first order in the Magnus expansion [38], the effective unitary operator is

$$U(t) \simeq \exp\left[\frac{-i}{2} \int_0^\tau ds \eta(t) \{\cos[\chi(s)] \sigma_z + \sin[\chi(s)] \sigma_y\}\right],$$

where $\chi(t) = \int_0^t c(s) ds$. The fidelity is then

$$\begin{aligned} \mathcal{F}(\tau) &= \left\langle \frac{1}{2} \cos^2\left(\int_0^\tau ds \eta(t) \{\cos[\chi(s)] \sigma_z \right. \right. \\ &\quad \left. \left. + \sin[\chi(s)] \sigma_y\}\right) \right\rangle_\eta \\ &= \frac{1}{2} + \frac{1}{2} \exp\left(-\frac{1}{2} \int_0^\tau ds_1 \int_0^{s_1} ds_2 \langle \eta(s_1) \eta(s_2) \rangle_\eta \right. \\ &\quad \left. \times \cos[\chi(s_1) - \chi(s_2)]\right) \\ &= \frac{1}{2} + \frac{1}{2} \exp\left(-\int_{-\infty}^\infty d\omega S(\omega) F(\omega)\right), \end{aligned}$$

where the filter function is defined as

$$\begin{aligned} F(\omega; \tau) &= \frac{1}{\omega^2} \int_0^\tau ds_1 \int_0^\tau ds_2 \cos[\chi(s_1) - \chi(s_2)] \\ &\quad \times \cos[(s_1 - s_2)\omega]. \end{aligned}$$

For EGP, $\chi(t) = \omega t$, while for dynamical decoupling, $\chi(t) = n_t \pi$, where n_t is the number of π pulses applied up to time t . The resulting filter functions are shown in Fig. 2. Note that for dynamical decoupling with hard pulses, the effective Hamiltonian converges at first order in the Magnus expansion. Higher-order terms will appear for the continuous driving case but will be negligible for weak noise and short time.

-
- [1] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser, *Quantum Computation by Adiabatic Evolution*, [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
 - [2] A.M. Childs, E. Farhi, and J. Preskill, *Robustness of Adiabatic Quantum Computation*, *Phys. Rev. A* **65**, 012322 (2001).
 - [3] F. Gaitan, *Simulation of Quantum Adiabatic Search in the Presence of Noise*, *Int. J. Quantum. Inform.* **04**, 843 (2006).
 - [4] M. Tiersch and R. Schützhold, *Non-Markovian Decoherence in the Adiabatic Quantum Search Algorithm*, *Phys. Rev. A* **75**, 062313 (2007).
 - [5] M.H.S. Amin, C.J.S. Truncik, and D.V. Averin, *Role of Single-Qubit Decoherence Time in Adiabatic Quantum Computation*, *Phys. Rev. A* **80**, 022303 (2009).
 - [6] I. de Vega, M.C. Bañuls, and A. Pérez, *Effects of Dissipation on an Adiabatic Quantum Search Algorithm*, *New J. Phys.* **12**, 123010 (2010).
 - [7] Stephen P. Jordan, Edward Farhi, and Peter W. Shor, *Error-Correcting Codes for Adiabatic Quantum Computation*, *Phys. Rev. A* **74**, 052322 (2006).
 - [8] D.A. Lidar, *Towards Fault Tolerant Adiabatic Quantum Computation*, *Phys. Rev. Lett.* **100**, 160506 (2008).
 - [9] Gerardo A. Paz-Silva, A. T. Rezakhani, Jason M. Dominy, and D.A. Lidar, *Zeno Effect for Quantum Computation and Control*, *Phys. Rev. Lett.* **108**, 080501 (2012).
 - [10] P. Facchi, D.A. Lidar, and S. Pascazio, *Unification of Dynamical Decoupling and the Quantum Zeno Effect*, *Phys. Rev. A* **69**, 032314 (2004).
 - [11] Mohan Sarovar and Kevin C. Young, *Error Suppression and Error Correction in Adiabatic Quantum Computation II: Non-equilibrium Dynamics*, [arXiv:1307.5892](https://arxiv.org/abs/1307.5892).
 - [12] Jérémie Roland and Nicolas J. Cerf, *Quantum Search by Local Adiabatic Evolution*, *Phys. Rev. A* **65**, 042308 (2002).
 - [13] A. T. Rezakhani, W.-J. Kuo, A. Hamma, D. A. Lidar, and P. Zanardi, *Quantum Adiabatic Brachistochrone*, *Phys. Rev. Lett.* **103**, 080502 (2009).
 - [14] Kevin C. Young, Robin Blume-Kohout, and Daniel Lidar, *Adiabatic Quantum Optimization with the Wrong Hamiltonian*, [arXiv:1310.0529](https://arxiv.org/abs/1310.0529).
 - [15] Daniel Gottesman, Ph.D. thesis, California Institute of Technology, 1997.
 - [16] Frank Gaitan, *Quantum Error Correction and Fault Tolerant Quantum Computing* (CRC Press, Boca Raton, FL, 2008).
 - [17] We define the weight of the Pauli operator to be the number of nonidentity terms in the tensor product. So,

- for example, $XXIII Z$ has weight three, denoted as $w(XXIII Z) = 3$.
- [18] Lorenza Viola, Emanuel Knill, and Seth Lloyd, *Dynamical Decoupling of Open Quantum Systems*, *Phys. Rev. Lett.* **82**, 2417 (1999).
- [19] Gregory Quiroz and Daniel A. Lidar, *High-Fidelity Adiabatic Quantum Computation via Dynamical Decoupling*, *Phys. Rev. A* **86**, 042333 (2012).
- [20] Götz S. Uhrig, *Keeping a Quantum Bit Alive by Optimized π -Pulse Sequences*, *Phys. Rev. Lett.* **98**, 100504 (2007).
- [21] N. Cody Jones, Thaddeus D. Ladd, and Bryan H. Fong, *Dynamical Decoupling of a Qubit with Always-On Control Fields*, *New J. Phys.* **14**, 093045 (2012).
- [22] Anand Ganti, Uzoma Onunkwo, and Kevin C. Young, *A Family of $[[6k, 2k, 2]]$ Codes for Practical, Scalable Adiabatic Quantum Computation*, [arXiv:1309.1674](https://arxiv.org/abs/1309.1674).
- [23] Graeme Smith and John A. Smolin, *Degenerate Quantum Codes for Pauli Channels*, *Phys. Rev. Lett.* **98**, 030501 (2007).
- [24] Generalization to degenerate codes simply requires us to interpret the operators E_{ν} as the correction operators associated with syndrome ν rather than with errors.
- [25] Os Vy, Xiaoting Wang, and Kurt Jacobs, *Error-Transparent Evolution: The Ability of Multi-body Interactions to Bypass Decoherence*, *New J. Phys.* **15**, 053002 (2013).
- [26] A. Yu. Kitaev, *Fault-Tolerant Quantum Computation by Anyons*, *Ann. Phys. (Amsterdam)* **303**, 2 (2003).
- [27] Alastair Kay, *Nonequilibrium Reliability of Quantum Memories*, *Phys. Rev. Lett.* **102**, 070503 (2009).
- [28] Sergey Bravyi and Barbara Terhal, *A No-Go Theorem for a Two-Dimensional Self-Correcting Quantum Memory Based on Stabilizer Codes*, *New J. Phys.* **11**, 043029 (2009).
- [29] Olivier Landon-Cardinal and David Poulin, *Local Topological Order Inhibits Thermal Stability in 2D*, *Phys. Rev. Lett.* **110**, 090502 (2013).
- [30] Jeongwan Haah, *Local Stabilizer Codes in Three Dimensions without String Logical Operators*, *Phys. Rev. A* **83**, 042330 (2011).
- [31] Sergey Bravyi and Jeongwan Haah, *Energy Landscape of 3D Spin Hamiltonians with Topological Order*, *Phys. Rev. Lett.* **107**, 150504 (2011).
- [32] Sergey Bravyi and Jeongwan Haah, *Analytic and Numerical Demonstration of Quantum Self-Correction in the 3D Cubic Code*, [arXiv:1112.3252](https://arxiv.org/abs/1112.3252).
- [33] Kamil Michnicki, *3-D Quantum Stabilizer Codes with a Power Law Energy Barrier*, [arXiv:1208.3496](https://arxiv.org/abs/1208.3496).
- [34] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill, *Topological Quantum Memory*, *J. Math. Phys. (N.Y.)* **43**, 4452 (2002).
- [35] M. W. Johnson, P. Bunyk, F. Maibaum, E. Tolkacheva, A. J. Berkley, E. M. Chapple, R. Harris, J. Johansson, T. Lanting, I. Perminov, E. Ladizinsky, T. Oh, and G. Rose, *A Scalable Control System for a Superconducting Adiabatic Quantum Optimization Processor*, *Supercond. Sci. Technol.* **23**, 065004 (2010).
- [36] Stephen P. Jordan and Edward Farhi, *Perturbative Gadgets at Arbitrary Orders*, *Phys. Rev. A* **77**, 062329 (2008).
- [37] For example, in gases, the dominant thermalization mechanism is three-body collisions.
- [38] J. A. Oteo and J. Ros, *From Time-Ordered Products to Magnus Expansion*, *J. Math. Phys. (N.Y.)* **41**, 3268 (2000).