

Fault-tolerant quantum computation versus Gaussian noise

Hui Khoon Ng and John Preskill

Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA

(Received 28 October 2008; published 16 March 2009)

We study the robustness of a fault-tolerant quantum computer subject to Gaussian non-Markovian quantum noise, and we show that scalable quantum computation is possible if the noise power spectrum satisfies an appropriate “threshold condition.” Our condition is less sensitive to very-high-frequency noise than previously derived threshold conditions for non-Markovian noise.

DOI: [10.1103/PhysRevA.79.032318](https://doi.org/10.1103/PhysRevA.79.032318)

PACS number(s): 03.67.Pp, 03.67.Lx

I. INTRODUCTION

The theory of fault-tolerant quantum computation shows that properly encoded quantum information can be protected against decoherence and processed reliably with imperfect hardware [1]. Demonstrating that this theory really works in practice is one of the great challenges facing contemporary science. A large-scale fault-tolerant quantum computer would be a scientific milestone, and it should also be *useful*, capable of solving hard problems that are beyond the reach of ordinary digital computers.

Though the theory of quantum fault tolerance strengthens our confidence that truly scalable quantum computers can be realized in the next few decades, failure is certainly possible. Perhaps the engineering challenges will prove to be so daunting, and the resources needed to overcome them so demanding, that society will be unable or unwilling to bear the cost for the foreseeable future. Perhaps new fundamental principles of physics, as yet undiscovered, will prevent large-scale quantum computers from behaving as currently accepted theory dictates. Finding that quantum computers fail for a fundamental reason would be a significant scientific advance but would disappoint prospective users.

There is a third reason to worry about the future prospects for fault-tolerant quantum computing. Mathematical results establishing that fault tolerance works effectively are premised on assumptions about the properties of the noise. The most obvious requirement is that the noise must be sufficiently weak—if the noise strength is below a *threshold of accuracy* then quantum computing is scalable in principle. But in addition, the noise must be suitably local, both spatially and temporally. Perhaps the quest for a quantum computer will be frustrated because the noise afflicting actual hardware is just not amenable to fault-tolerant protocols.

We can anticipate therefore that progress toward scalable quantum computing will require an ongoing dialog between experimenters who will better understand the limitations of their devices and theorists who will propose better ways to overcome the limitations and to evaluate the efficacy of these proposals. In the meantime, an important task for theorists is to broaden the range of noise models for which useful accuracy threshold theorems can be proven, and we pursue that task in this paper. Our main result is a proof of the threshold theorem for non-Markovian Gaussian noise models, in which system qubits are locally coupled to bath variables that have Gaussian fluctuations. Specifically, if the bath is a system of

uncoupled harmonic oscillators, at either zero or nonzero temperature, our theorem expresses the threshold condition in terms of the power spectrum of the bath fluctuations.

Early proofs of the threshold theorem [2–4] assumed that the noise is *Markovian*. This means that each quantum gate in the noisy circuit can be modeled as a unitary transformation that acts jointly on a set of the qubits in the computer (the system qubits) and on the environment (the bath variables), but where it is assumed that the bath has no memory—the state of the bath is refreshed after every gate. The theorem was extended to a class of non-Markovian noise models in [5] and further generalized in [6,7]. The results in [6,7] have the substantial virtue that the state of the bath and its internal dynamics can be arbitrary; for fault-tolerant quantum computing to work, it is only required that the bath couples weakly and locally to the system.

However these results also have two serious drawbacks. First, the threshold condition is not easily related to experimentally accessible quantities, rather it requires terms in the Hamiltonian that couple the system to the bath to have a sufficiently small operator norm. Second, this condition severely constrains the very-high-frequency fluctuations of the bath. Intuitively, it seems that this constraint, which may limit the applicability of the threshold theorem to noise in some realistic settings, ought not to be necessary since fluctuations with a time scale much shorter than the time it takes to execute a quantum gate tend to average out.

One possible way to reach more pleasing conclusions is to make physically reasonable assumptions about the noise that go beyond the assumptions in [6,7]; that is the approach we follow here. Our threshold theorem applies to any noise model in which the bath variables are free fields (aside from their coupling to the system qubits) and expresses the threshold condition in terms of the bath’s two-point correlation function, which is in principle measurable. It should be possible to extend our analysis to the case where the bath variables have sufficiently weak self-interactions, though we will not pursue that extension here. Furthermore, though our threshold condition still requires the very-high-frequency bath fluctuations to be sufficiently weak, this requirement is considerably relaxed compared to previous threshold theorems that apply to non-Markovian noise. Showing that these requirements can be relaxed even further, perhaps by making additional physically motivated assumptions, is an important open problem.

Experimenters use a variety of techniques to suppress the noise in quantum hardware, such as cleverly designed pulse

sequences to improve the fidelity of quantum gates (spin echos, dynamical decoupling, etc.) and intrinsically robust encodings of quantum information (noiseless subsystems, topologically protected qubits, etc.). These techniques can be highly effective and are likely to be incorporated into the design of future quantum computers, but do not by themselves suffice to ensure the scalability of quantum computing. After such tricks are exhausted some residual noise inevitably remains that must be controlled using quantum error-correcting codes and fault-tolerant methods. Since our objective in this paper is to study the effectiveness of these fault-tolerant methods, our noise models may be viewed as (perhaps highly idealized) effective descriptions of this residual noise in “fundamental” quantum gates that might already be realized using complex and sophisticated protocols.

After reviewing previously known formulations of the quantum accuracy threshold theorem in Sec. II (with some details relegated to the Appendix), we state our result in Sec. III, explore some of its implications in Sec. IV, derive it in Sec. V, and discuss some generalizations in Sec. VI. We derive a sharper result for the case of pure dephasing noise in Sec. VII. Section VIII contains our conclusions.

II. NOISE MODELS AND QUANTUM ACCURACY THRESHOLD THEOREMS

Here we will briefly review some previously known formulations of the quantum accuracy threshold theorem and explain why these results still leave something to be desired. Then in Sec. III we will state our result, which addresses some of the shortcomings of the previous results.

The goal of fault-tolerant quantum computing is to simulate an ideal quantum circuit using the noisy gates that can be executed by actual devices. Theoretical results show that this goal is attainable if the noise is not too strong and not too strongly correlated. The essential trick that makes fault tolerance work is that the logical quantum state processed by the computer can be encoded very nonlocally, so that it is well protected from damage caused by local noise.

It is convenient to analyze the effectiveness of a fault-tolerant noisy circuit by invoking a *fault-path expansion*, schematically,

$$\text{noisy circuit} = \sum \text{“fault path.”} \quad (1)$$

Let us use the term *location* to speak of an operation in a quantum circuit that is performed in a single time step; a location may be a single-qubit or multiqubit gate, a qubit preparation step, a qubit measurement, or the identity operation in the case of a qubit that is idle during the time step. In each fault path, the quantum gates are faulty at a specified set of locations in the circuit, while at all other locations the quantum gates are assumed to be ideal. We say that the faulty locations are “bad” and that the ideal locations are “good.” The general concept of a fault-path expansion applies quite broadly, and different noise models can be distinguished according to how we flesh out the meaning of Eq. (1).

A. Local stochastic noise

In a “stochastic” noise model we assign a *probability* to each fault path [6]. We speak of *local stochastic noise* with strength ϵ if, for any specified set \mathcal{I}_r of r locations in the circuit, the sum $P^{\text{bad}}(\mathcal{I}_r)$ of the probabilities of all fault paths that are bad at *all* of these r locations satisfies

$$P^{\text{bad}}(\mathcal{I}_r) \leq \epsilon^r. \quad (2)$$

In this noise model, no further restrictions are imposed on the noise, and in particular the trace-preserving quantum operation applied at the faulty locations may be chosen for each fault path by an adversary who wants the computation to fail. Thus the faults can be correlated, both spatially and temporally, but the adversary’s power is limited because an attack on r specified circuit locations occurs with probability at most ϵ^r . The noise is “local” in the sense that attacking each additional location suppresses the probability of the fault path by another power of ϵ .

Most proofs of the threshold theorem use *recursive simulations*. This means that quantum information is protected by a hierarchy of codes within codes and that the fault-tolerant circuit has a self-similar structure. We refer to an unencoded quantum circuit as a “level-0” simulation. In a level-1 simulation, each elementary gate in the level-0 circuit is replaced by a level-1 *gadget* constructed from elementary gates; this 1-gadget performs the appropriate encoded operation on logical qubits that are protected by a quantum error-correcting code \mathcal{C} . In a level-2 simulation, each elementary gate in the ideal circuit is replaced by a level-2 gadget; the 2-gadget is constructed by replacing each elementary gate in the 1-gadget by a 1-gadget. A 2-gadget operates on quantum information protected by $\mathcal{C} \triangleright \mathcal{C}$, where \triangleright denotes code concatenation. (That is, $\mathcal{C}_1 \triangleright \mathcal{C}_2$ is encoded by first encoding the “outer” code \mathcal{C}_2 and then encoding each qubit in the \mathcal{C}_2 block using the “inner” code \mathcal{C}_1 .) In a level- k simulation, each elementary gate in the ideal circuit is replaced by a level- k gadget, constructed by replacing each elementary gate in the $(k-1)$ -gadget by a 1-gadget; it operates on quantum information protected by $\mathcal{C}^{\triangleright k}$.

For local stochastic noise and also for other noise models with suitable properties, a recursive simulation can be analyzed by a procedure called *level reduction*, in which a level- k simulation is mapped to a “coarse-grained” level- $(k-1)$ simulation that acts on the top-level logical information in exactly the same way. Suppose, for example, that \mathcal{C} is a distance-3 code that can correct one error. Then if the 1-gadgets are properly designed, each good 1-gadget that contains no more than one faulty location simulates the corresponding ideal gate correctly, while bad 1-gadgets with more than one fault may simulate the ideal gate incorrectly. In the level-reduction step, for each fault path the good 1-gadgets are mapped to ideal level-0 gates, while the bad 1-gadgets are mapped to faulty level-0 gates. After this step, the resulting noisy circuit is still subject to local stochastic noise but with a renormalized value of the noise strength,

$$\epsilon^{(1)} = \epsilon^2 / \epsilon_0 = \epsilon_0 (\epsilon / \epsilon_0)^2. \quad (3)$$

The renormalized value of the noise strength is $O(\epsilon^2)$ because at least two faults are required for a 1-gadget to fail;

the quantity ε_0^{-1} is a combinatoric factor counting the number of “malignant” sets of locations within the 1-gadget where faults can cause failure.

Since level reduction maps local stochastic noise to local stochastic noise (but with a revised value of the noise strength), the level-reduction step can be carried out repeatedly and analyzed by the same method each time. That the structure of the noise is preserved, even though its strength is renormalized, is a useful feature of the local stochastic noise model not shared by some noise models. For example, if faults in level-0 gates were independently and identically distributed, the effective noise model after one level-reduction step would become correlated rather than independent. See [6,8] for a more detailed discussion of the level-reduction procedure.

By repeating the level-reduction step all together k times, we reduce the level- k simulation to an effective level-0 (i.e., unencoded) simulation with noise strength,

$$\varepsilon^{(k)} = \varepsilon_0(\varepsilon/\varepsilon_0)^{2^k}. \quad (4)$$

It follows that for $\varepsilon < \varepsilon_0$ (the *accuracy threshold*), the effective noise strength becomes negligibly small for k sufficiently large, and the simulation becomes highly reliable. More precisely, for any fixed $\varepsilon < \varepsilon_0$ and fixed $\delta > 0$, an ideal circuit with L gates can be simulated with error probability δ by a noisy circuit with L^* gates, where for some constant c

$$L^*/L = O\left(\left(\frac{\ln(L/\delta)}{\ln(\varepsilon_0/\varepsilon)}\right)^c\right). \quad (5)$$

(The constant c is determined by the size of the 1-gadgets.) Thus, with reasonable overhead cost, the noisy simulation gets the right answer with high probability. This is the quantum accuracy threshold theorem for local stochastic noise.

For the threshold theorem to apply, two features of the simulation are essential. First, we must assume that quantum gates can be executed in parallel—otherwise we would be unable to control storage errors that occur simultaneously in different parts of the computer. Second, we assume that qubits can be “discarded” and replaced by fresh qubits (for example, by measuring the qubits and resetting them)—otherwise we would be unable to flush from the computer the entropy introduced by noise. Estimates of the accuracy threshold ε_0 often rely on further assumptions. For example, if we assume that qubit measurements are as fast as quantum gates, that classical computations are arbitrarily accurate, that the accuracy of a two-qubit quantum gate does not depend on the spatial separation of the qubits, and that no data qubits “leak” from the computational Hilbert space, then it has been shown that $\varepsilon_0 > 0.67 \times 10^{-3}$ [9]. For noise models with weaker correlations than in the local stochastic noise model, the proven accuracy threshold is above 10^{-3} [10,11], and numerical evidence suggests that the actual value of the threshold can be of order 1% [12,13]. Furthermore, it has been shown that the threshold is not drastically reduced if some of these assumptions are relaxed, for example, by allowing measurements to be slow [14], allowing leakage [15], or requiring quantum gates to be local on a two-dimensional array [16]. Upper bounds on the accuracy threshold have also

been derived. Recent estimates, which apply to restricted types of fault-tolerant architectures, indicate that ε_0 cannot be much larger than 5% [17,18].

B. Local non-Markovian noise

The local stochastic noise model is handy for analysis and has some quasirealistic features, but it is still rather artificial. From a physics perspective, it is more natural to formulate the noise model in terms of a Hamiltonian H that governs the joint evolution of the system and the bath. We may express H as

$$H = H_S + H_B + H_{SB}, \quad (6)$$

where H_S is the time-dependent system Hamiltonian that realizes the ideal quantum circuit, H_B is the (arbitrary) Hamiltonian of the bath, and H_{SB} is a perturbation, responsible for the noise, which may couple the system to the bath. For this continuous-time description of a quantum circuit, a location consists of a specified qubit or set of qubits to which a gate is applied and a specified *time interval* during which that gate is realized by the ideal system Hamiltonian H_S .

We say that such a noise model is *non-Markovian*, meaning that quantum information can escape from the system to the bath and then return to the system at a later time, so that the state of the system at time $t+dt$ is not uniquely determined by its state at time t . Furthermore, H_{SB} may also contain terms that act nontrivially only on the system, representing unitary noise arising from imperfect control of the system Hamiltonian. Actually, the local stochastic noise model already incorporates some non-Markovian effects; even when fault paths are weighted by probabilities, the adversary who attacks the circuit might employ a quantum memory. But different methods are needed to analyze the consequences of Hamiltonian noise models because fault paths are summed coherently rather than stochastically.

The locations in a quantum circuit include not only quantum gates and storage steps but also qubit preparation and measurement steps. Preparation and measurement noise can be incorporated into a Hamiltonian description by various means. In this paper we will take an especially simple approach, modeling an imperfect preparation by an ideal preparation followed by evolution governed by H and modeling an imperfect measurement by an ideal measurement preceded by evolution governed by H . For the time being, to simplify the discussion, we will imagine that system qubits are prepared only at the very beginning of the computation and measured only at the very end. Preparations and measurements that occur at intermediate times can easily be incorporated. Indeed, including such preparations and measurements is essential because quantum computing is scalable only if qubits can be regularly refreshed during the course of the computation. We will elaborate on this point in Sec. VI.

We may say that the Hamiltonian noise model is local if the perturbation H_{SB} can be expressed as a sum of terms,

$$H_{SB} = \sum_a H_{SB}^{(a)}, \quad (7)$$

where each $H_{SB}^{(a)}$ acts on only a small number of system qubits (while perhaps also acting collectively on many bath

variables). The joint unitary time-evolution operator U_{SB} for system and bath, resulting from integrating the Schrödinger equation for Hamiltonian H , can be formally expanded to all orders in time-dependent perturbation theory in H_{SB} . In any fixed term in this expansion, perturbations chosen from the set $\{H_{SB}^{(a)}\}$ are inserted at specified times. For such a fixed term in the perturbation expansion, let us say that a location in the (level-0) noisy simulation is bad if an inserted perturbation acts nontrivially somewhere inside that location; otherwise that location is good. Of course, under this definition a single insertion of $H_{SB}^{(a)}$ might cause two (or perhaps more) locations to be bad in a particular time step, if $H_{SB}^{(a)}$ acts collectively on two qubits that are undergoing different gates executed in parallel in the ideal circuit.

As already noted, we may assume that the system qubits have been initialized ideally at the start of the Hamiltonian evolution; we denote this initial system state by $|\Psi_S^0\rangle$. We also assume that the initial state of the bath is a pure state $|\Psi_B^0\rangle$. There is really no loss of generality in supposing that the bath starts out in a pure state; if we wish to consider a mixed initial state of the bath instead (for example, a thermal state), we may include in the bath a “reference” system that “purifies” the mixed state.

We have also noted that we may assume that final measurements performed on system qubits are ideal. Just before these final measurements are conducted, the (pure) state of the system and bath is

$$|\Psi_{SB}\rangle = U_{SB}|\Psi_{SB}^0\rangle, \tag{8}$$

where $|\Psi_{SB}^0\rangle = |\Psi_S^0\rangle \otimes |\Psi_B^0\rangle$ is the initial state of system and bath. For any specified set \mathcal{I}_r of r locations in the circuit, let us denote by $|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r)\rangle$ the sum of all the terms in the formal perturbation expansion of $|\Psi_{SB}\rangle$ such that all of these r locations are bad. Then we speak of *local non-Markovian noise* (or more briefly *local noise*) with strength ε if

$$\| |\Psi_{SB}^{\text{bad}}(\mathcal{I}_r)\rangle \| \leq \varepsilon^r. \tag{9}$$

The noise strength ε can be related to properties of the perturbation H_{SB} . We will sometimes refer to this model as *local coherent* noise to emphasize that (in contrast to the local stochastic noise model) fault paths are assigned amplitudes rather than probabilities.

Although there are some new subtleties (see [6] and the Appendix), the level-reduction concept can be applied to Hamiltonian noise models in much the same way as for stochastic models. We may say that a 1-gadget is bad if it contains bad level-0 gates at a malignant set of locations, that a 2-gadget is bad if it contains bad 1-gadgets at a malignant set of locations, that a 3-gadget is bad if it contains bad 2-gadgets at a malignant set of locations, and so on. For any specified set $\mathcal{I}_r^{(k)}$ of r k -gadgets in the circuit, let us denote by $|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(k)})\rangle$ the sum of all the terms in the formal perturbation expansion of $|\Psi_{SB}\rangle$ such that all of these r k -gadgets are bad. Then it follows from Eq. (9) that

$$\| |\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(k)})\rangle \| \leq (\varepsilon^{(k)})^r, \tag{10}$$

with $\varepsilon^{(k)}$ as in Eq. (4); the derivation of Eq. (10) is sketched in the Appendix. Furthermore, a level- k simulation in which

all k -gadgets are good simulates the ideal circuit perfectly. In this sense, repeated level reduction reduces a level- k simulation to an equivalent level-0 simulation while mapping local noise to local noise with a renormalized noise strength $\varepsilon^{(k)}$, and for $\varepsilon < \varepsilon_0$, the renormalized noise strength becomes negligible for large k . The threshold value ε_0 of the noise strength for local noise is of the same order (though not exactly the same) as the threshold for local stochastic noise. We emphasize that, once Eq. (9) is established, we can derive Eq. (10) without any further assumptions about the Hamiltonian $H = H_S + H_B + H_{SB}$.

The strength ε of local noise can be estimated based on the detailed properties of the expansion in Eq. (7) of the perturbation H_{SB} in terms of local system operators. For example, in [6] the noise was assumed to be “short range” in the sense that the perturbation H_{SB} acts collectively on a pair of data qubits only while the ideal system Hamiltonian H_S also couples those two data qubits—that is, only while the ideal quantum circuit calls for that pair of qubits to undergo a two-qubit gate. For this short-range local noise model, it was shown that Eq. (9) is satisfied if we choose

$$\varepsilon = (\max_{a,t} \|H_{SB}^{(a)}(t)\|) t_0, \tag{11}$$

where t_0 is the time needed to execute a quantum gate, $\|\cdot\|$ denotes the sup operator norm, and the maximum is over all circuit locations and all times. On the other hand, in [7] the noise was assumed to be “long range” with H_{SB} coupling each pair of data qubits irrespective of the structure of the ideal circuit. In that case we may write

$$H_{SB} = \sum_{\langle ij \rangle} H_{\langle ij \rangle}, \tag{12}$$

where the sum is over all unordered pairs of system qubits; $H_{\langle ij \rangle}$ acts collectively on the pair of qubits $\langle ij \rangle$ and also on the bath. For this long-range local noise model, it was shown that Eq. (9) is satisfied if we choose

$$\varepsilon^2 = C \left(\max_{i,t} \sum_j \|H_{\langle ij \rangle}\| \right) t_0, \tag{13}$$

where C is the numerical constant $C = 2e \approx (2.34)^2$. (This is actually a slight improvement over the value of C reported in [7]; the improved value can be derived using the reasoning described in Sec. V C below, if we assume $\varepsilon^2 \leq e$.)

The origin of Eq. (11) is easy to understand intuitively [5]. If each one of the r specified locations in \mathcal{I}_r is bad, then the perturbation must be inserted at least once in each of these locations, and each insertion reduces the norm of the state by a factor of at least $\|H_{SB}^{(a)}\|$. Inside each location, there is an *earliest* insertion of the perturbation that can occur at any time during the duration of the location, a time window of width t_0 . Integrating over the time of the earliest insertion of the perturbation inside each location, we obtain Eq. (11). For the long-range noise model, a single insertion of the perturbation $H_{\langle ij \rangle}$ can cause two circuit locations to be bad if qubits i and j are participating in separate gates, and therefore the noise strength is correspondingly higher [observe that ε^2 rather than ε appears on the left-hand side of Eq. (13)].

C. Assessment

The results [Eqs. (11) and (13)] are significant because they demonstrate that quantum computing is scalable, in principle, for non-Markovian noise described by a system-bath Hamiltonian. Furthermore, this formulation of the threshold theorem has the noteworthy advantage that the argument works for any bath Hamiltonian H_B . The dynamics of the bath does not matter, as long as the perturbation H_{SB} is local and sufficiently weak.

However, expressing the threshold condition as in Eqs. (11) and (13) has serious drawbacks. First we should note that while in the local stochastic noise model we may interpret the noise strength ε as an error probability per gate, in the non-Markovian noise model ε is really an error *amplitude*. Since a probability is a square of an amplitude, requiring $\varepsilon < \varepsilon_0$ in the local noise model is a far more stringent criterion than requiring $\varepsilon < \varepsilon_0$ in the local stochastic noise model. Our analysis yields a much weaker lower bound on the accuracy threshold for the local noise model than for the local stochastic noise model because we pessimistically allow the bad fault paths to add together with a common phase and thus to interfere constructively.

Most likely this analysis is far too pessimistic. For one thing, two fault paths cannot interfere if they generate different error syndrome histories [6]. Furthermore, it is reasonable to expect that distinct fault paths with the same syndrome history have only weakly correlated phases, and if so, then the modulus of a sum of N fault paths should grow like \sqrt{N} rather than linearly in N . That is, if the phases of fault paths can be regarded as random, then we expect the *probabilities* of the fault paths, rather than their amplitudes, to accumulate linearly. An important open problem for the theory of quantum fault tolerance is to put this phase randomization hypothesis on a rigorous footing, which would establish a much higher estimate of the accuracy threshold for local noise. But we will not be addressing this problem in this paper.

There are other drawbacks of the threshold condition Eq. (11) that we *will* try to address, however. One issue is that the norm of the system-bath Hamiltonian is not directly measurable in experiments, and it would be far preferable to state the threshold condition in terms of experimentally accessible quantities, such as the noise power spectrum. In fact, for otherwise reasonable noise models, the norm $\|H_{SB}^{(a)}\|$ could be formally infinite (if, for example, the system qubits couple to unbounded bath operators such as the quadrature amplitudes of bath oscillators), and in such cases the threshold theorem has little force.

In more physical terms, an undesirable feature of demanding small ε where ε is given by Eq. (11) is that this condition requires that the very-high-frequency component of the noise be particularly weak, a requirement that seems not to be physically well motivated. To be concrete, suppose that

$$H_{SB}^{(a)} = S^{(a)} \otimes B^{(a)}, \quad (14)$$

where $S^{(a)}$ is a local Hermitian system operator with $\|S^{(a)}\| = 1$ and $B^{(a)}$ is a Hermitian bath operator. Then combining the condition $\varepsilon < \varepsilon_0$ with Eq. (11) implies in particular that

$$\langle \Psi_{SB}^0 | B^{(a)}(t) B^{(a)}(t') | \Psi_{SB}^0 \rangle = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \tilde{\Delta}(\omega) < \varepsilon_0^2 / t_0^2, \quad (15)$$

where $\tilde{\Delta}(\omega)$ is the Fourier transform of the bath's two-point correlation function, defined by

$$\langle \Psi_{SB}^0 | B^{(a)}(t) B^{(a)}(t') | \Psi_{SB}^0 \rangle = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-i\omega(t-t')} \tilde{\Delta}(\omega) \quad (16)$$

[$B^{(a)}(t)$ denotes the interaction-picture bath operator]. Suppose that the fluctuations of the bath variables are Ohmic (and at zero temperature); that is, linear in frequency at low (positive) frequency and exponentially decaying at frequencies large compared to the cutoff frequency τ_c^{-1} ,

$$\tilde{\Delta}(\omega) = \begin{cases} 2\pi A \omega e^{-\omega\tau_c} & \text{if } \omega \geq 0 \\ 0 & \text{if } \omega < 0, \end{cases} \quad (17)$$

where A is a positive dimensionless parameter quantifying the strength of the Ohmic noise. Then the threshold condition implies that

$$\sqrt{A} \left(\frac{t_0}{\tau_c} \right) < \varepsilon_0. \quad (18)$$

For the case of Ohmic noise, then, the quantity that is required to be small is linearly “ultraviolet divergent”; that is, it has a linear sensitivity to the high-frequency cutoff τ_c^{-1} , which may be orders of magnitude higher than the characteristic frequency τ_0^{-1} of the ideal computation.

The extreme sensitivity of the threshold condition to the very-high-frequency noise seems surprising since one's naive expectation is that noise with zero mean and frequency much larger than τ_0^{-1} should nearly average out. This unsatisfying limitation of Eq. (11), already pointed out in the original paper by Terhal and Burkard [5] (and later highlighted by Alicki [19] and by Hines and Stamp [20]), may just be a shortcoming of the analysis, but conceivably it hints at a deeper problem for quantum fault tolerance. For example, it has been suggested [21] that during the course of a long quantum computation, an initially benign state of the bath may be pushed toward a far more malicious state that compromises the fault-tolerant protocol. Perhaps high-frequency noise with zero mean, which locally seems incapable of inflicting serious harm, has cumulative global effects that are surprisingly troublesome. Whether or not one suspects that the environment could be so cunning an adversary, stronger rigorous arguments establishing that quantum computing is robust against non-Markovian noise would surely be welcome.

Our central result in this paper is an estimate of the noise strength ε that applies to a Hamiltonian description of Gaussian non-Markovian noise. We will formulate the noise model and state our result in Sec. III, discuss some implications in Sec. IV, and postpone the derivation until Sec. V. For this particular important class of noise models, we will be able to state a threshold condition that is less sensitive to very-high-frequency noise, though some sensitivity will still remain.

The combinatoric analysis that leads to our result borrows substantially from the derivation in [7] of Eq. (13), though the context is rather different.

III. GAUSSIAN NOISE AND THE THRESHOLD CONDITION

By ‘‘Gaussian noise’’ we mean a Hamiltonian noise model where the bath is a set of uncoupled harmonic oscillators, and each system qubit couples to a linear combination of oscillator quadrature amplitudes; hence (in units with $\hbar=1$)

$$H_B = \sum_k \omega_k a_k^\dagger a_k, \quad (19)$$

and

$$H_{SB} = \sum_x \sum_\alpha \sigma_\alpha(x) \otimes \tilde{\phi}_\alpha(x, t), \quad (20)$$

where $\tilde{\phi}_\alpha(x, t)$ is the Hermitian operator,

$$\tilde{\phi}_\alpha(x, t) = \sum_k [g_{k,\alpha}(x, t) a_k + g_{k,\alpha}^*(x, t) a_k^\dagger]. \quad (21)$$

Here x is a label indicating a system qubit’s position, and $\{\sigma_\alpha(x), \alpha=1, 2, 3\}$ are the three Pauli operators acting on qubit x . The a_k ’s are annihilation operators for the bath oscillators, satisfying the commutation relation $[a_k, a_{k'}^\dagger] = \delta_{kk'}$, and $g_{k,\alpha}(x, t)$ is a complex coupling parameter that determines how strongly oscillator k couples to qubit x at time t .

As explained in Sec. II B (see also Sec. VI A), we may assume without loss of generality that the bath is prepared in a pure state $|\Psi_B^0\rangle$ at the beginning of the computation. The Hamiltonian $H_B + H_{SB}$, along with the choice of the bath’s initial state $|\Psi_B^0\rangle$, defines our noise model. The bath fluctuations will be Gaussian if the state $|\Psi_B^0\rangle$ is a Gaussian state (that is, a generalized ‘‘squeezed’’ state) of the oscillator bath—a purified thermal state is a special case of such a squeezed state.

It is useful to define the ‘‘interaction-picture’’ bath operator $\phi_\alpha(x, t)$ as

$$\begin{aligned} \phi_\alpha(x, t) &= e^{iH_B t} \tilde{\phi}_\alpha(x, t) e^{-iH_B t} \\ &= \sum_k [g_{k,\alpha}(x, t) a_k e^{-i\omega_k t} + g_{k,\alpha}^*(x, t) a_k^\dagger e^{i\omega_k t}], \end{aligned} \quad (22)$$

and to define the bath’s two-point correlation function as

$$\Delta(\alpha_1, x_1, t_1; \alpha_2, x_2, t_2) = \langle \Psi_B^0 | \phi_{\alpha_1}(x_1, t_1) \phi_{\alpha_2}(x_2, t_2) | \Psi_B^0 \rangle. \quad (23)$$

We will sometimes use the abbreviated notation $\phi(1)$ for $\phi_{\alpha_1}(x_1, t_1)$ and $\Delta(1, 2)$ for $\Delta(\alpha_1, x_1, t_1; \alpha_2, x_2, t_2)$; we also define

$$|\bar{\Delta}(1, 2)| = \sum_{\alpha_1, \alpha_2} |\Delta(\alpha_1, x_1, t_1; \alpha_2, x_2, t_2)|. \quad (24)$$

When we say that the noise is Gaussian, we mean that the bath variable $\phi_\alpha(x, t)$ obeys Gaussian statistics: all n -point bath correlation functions vanish for n odd, and the $2n$ -point

$$\langle \phi(1)\phi(2)\phi(3)\phi(4) \rangle = \text{diagram 1} + \text{diagram 2} + \text{diagram 3}$$

FIG. 1. (Color online) The four-point correlation function for a free field can be expressed in terms of products of two-point correlation functions by summing over all ‘‘contractions,’’ where each contraction divides the four points into two unordered pairs.

function can be expressed in terms of two-point functions. Using $\langle \cdot \rangle$ to denote the expectation value in the state $|\Psi_B^0\rangle$, Gaussian statistics implies that

$$\begin{aligned} \langle \phi(1)\phi(2)\phi(3)\cdots\phi(2n) \rangle \\ = \sum_{\text{contractions}} \Delta(i_1, i_2) \Delta(i_3, i_4) \cdots \Delta(i_{2n-1}, i_{2n}), \end{aligned} \quad (25)$$

where summing over ‘‘contractions’’ means summing over the $(2n)!/2^n n!$ ways to divide the labels $1, 2, 3, \dots, 2n$ into n unordered pairs. For example, if ϕ is a Gaussian variable, then the four-point function is

$$\begin{aligned} \langle \phi(1)\phi(2)\phi(3)\phi(4) \rangle &= \Delta(1, 2)\Delta(3, 4) + \Delta(1, 3)\Delta(2, 4) \\ &\quad + \Delta(1, 4)\Delta(2, 3), \end{aligned} \quad (26)$$

as illustrated in Fig. 1. This expansion of the $2n$ -point function in terms of two-point functions is sometimes called ‘‘Wick’s theorem.’’

Now we can state our main result: Gaussian noise obeys the local noise condition Eq. (9), with noise strength

$$\varepsilon^2 = C \max_{\text{loc}} \left(\int_{1, \text{loc}} \int_{2, \text{all}} |\bar{\Delta}(1, 2)| \right), \quad (27)$$

where $C=2e \approx (2.34)^2$ is the numerical constant defined earlier (and where we have assumed $\varepsilon^2 \leq e$). Here $\int_{1, \text{loc}}$ indicates that one leg (x_1, t_1) of the two-point function is integrated over a single location in the circuit: x_1 is summed over the qubits participating in a particular gate and t_1 is integrated over the time interval in which that gate is executed. $\int_{2, \text{all}}$ indicates that the other leg (x_2, t_2) of the two-point function is summed over all system qubits and integrated over the entire duration of the computation. The maximum is with respect to all possible circuit locations for (x_1, t_1) . The threshold condition $\varepsilon < \varepsilon_0$, with ε given by Eq. (27), now becomes a condition on the two-point correlation function of the bath. We note that the ordering of the operators $\phi(1)$ and $\phi(2)$ does not matter in Eq. (27) because $|\Delta(1, 2)| = |\Delta(2, 1)|$; changing the ordering modifies only the phase of $\Delta(1, 2)$, not its modulus.

Another noteworthy feature is that our estimate of ε applies for an arbitrary system Hamiltonian. This property may seem unexpected at first, as we know that in some settings the damage caused by the noise can depend on the relation between the energy spectrum of the H_S and the power spectrum of the noise. For example, the spontaneous decay rate for a qubit with energy splitting $\hbar\omega$ depends on the noise power at circular frequency ω . How, then, can our threshold condition depend only on the noise spectrum and not on the energy spectrum of H_S ? The answer is that by taking the modulus $|\bar{\Delta}(1, 2)|$ of the bath two-point function in Eq. (27)

we are already being maximally pessimistic about how the spectrum of H_S matches the noise power spectrum. Thus there are both advantages and disadvantages in formulating a threshold condition that is general enough to apply for any ideal system Hamiltonian. On the one hand we find a criterion for scalable quantum computing that can be stated easily and rigorously proved by a reasonably simple argument. On the other hand, the price of such rigor is that our stated criterion may be far more demanding than it really needs to be.

The crucial assumption in the derivation of Eq. (27) is Eq. (20), where $\tilde{\phi}_\alpha(x, t)$ is a “free field,” i.e., obeys Gaussian statistics; thus Eq. (21) could be regarded as merely a general phenomenological representation of a Gaussian field and not necessarily as a fundamentally accurate microscopic description of the bath. Caldeira and Leggett [22] argued that noise is expected to be Gaussian, at least to an excellent approximation, in a wide variety of realistic physical settings where the system is weakly coupled to many environmental degrees of freedom.

If the initial state of the bath is (the purification of) a thermal state with inverse temperature $\beta=1/kT$, then the mean occupation number of each oscillator is determined by the Bose-Einstein distribution function; we have

$$\begin{aligned} \langle \Psi_B^0 | a_k^\dagger a_{k'} | \Psi_B^0 \rangle &= \frac{\delta_{kk'}}{e^{\beta\omega_k} - 1} = \langle \Psi_B^0 | a_k a_{k'}^\dagger | \Psi_B^0 \rangle - 1, \\ \langle \Psi_B^0 | a_k a_{k'} | \Psi_B^0 \rangle &= 0 = \langle \Psi_B^0 | a_k^\dagger a_{k'}^\dagger | \Psi_B^0 \rangle, \end{aligned} \quad (28)$$

and therefore

$$\begin{aligned} \Delta(\alpha_1, x_1, t_1; \alpha_2, x_2, t_2) &= \frac{1}{2} \sum_k g_{k, \alpha_1}(x_1, t_1) g_{k, \alpha_2}^*(x_2, t_2) e^{-i\omega_k(t_1 - t_2)} \\ &\quad \times [\coth(\beta\omega_k/2) + 1] \\ &\quad + \frac{1}{2} \sum_k g_{k, \alpha_1}^*(x_1, t_1) g_{k, \alpha_2}(x_2, t_2) e^{i\omega_k(t_1 - t_2)} \\ &\quad \times [\coth(\beta\omega_k/2) - 1]. \end{aligned} \quad (29)$$

Just to be concrete, consider the case where the noise is stationary and spatially uncorrelated—each qubit has a time-independent coupling to its own independent oscillator bath (though admittedly these are dubious assumptions when multiqubit gates are executed). Then

$$\Delta(\alpha_1, x_1, t_1; \alpha_2, x_2, t_2) = \delta_{x_1 x_2} \Delta(\alpha_1, x_1, t_1; \alpha_2, x_1, t_2), \quad (30)$$

where

$$\Delta(\alpha_1, x_1, t_1; \alpha_2, x_1, t_2) = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-i\omega(t_1 - t_2)} \tilde{\Delta}_{\alpha_1 \alpha_2}(x_1, \omega) \quad (31)$$

and

$$\tilde{\Delta}_{\alpha_1 \alpha_2}(x_1, \omega) = \begin{cases} \pi J_{\alpha_1, \alpha_2}(x_1, \omega) [\coth(\beta\omega/2) + 1] & \text{if } \omega > 0 \\ \pi J_{\alpha_1, \alpha_2}^*(x_1, \omega) [\coth(\beta\omega/2) - 1] & \text{if } \omega < 0. \end{cases} \quad (32)$$

Here $J_{\alpha_1, \alpha_2}(x_1, \omega)$ is the Hermitian matrix,

$$J_{\alpha_1, \alpha_2}(x_1, \omega) = \sum_k \delta(\omega - \omega_k) g_{k, \alpha_1}(x_1) g_{k, \alpha_2}^*(x_1). \quad (33)$$

The function $J_{\alpha_1, \alpha_2}(x_1, \omega)$ is the spin-polarization-dependent power spectrum of the noise acting on qubit x_1 . If the energy splitting $\hbar\omega$ of the qubit is tunable, this function can be measured by observing the qubit’s relaxation rate as a function of the energy splitting and the polarization. In principle, n -qubit correlations in the noise can also be measured using quantum process tomography, if n is not too large.

IV. SOME IMPLICATIONS

Before presenting our derivation of Eq. (27) in Sec. V, we will discuss a few of its implications.

A. Dimensional criterion

Our expression for ε in Eq. (27) involves a formal integration over all space and time. If the bath correlations decay slowly in space or time, this integral might diverge in the limit of a computation that is very wide, very deep, or both. In that case, our “threshold condition” cannot be satisfied asymptotically, and we cannot conclude that quantum computation is scalable. On the other hand, if the integral converges “in the infrared,” then the threshold condition has value, as it establishes scalability if the coupling of the system to the bath is sufficiently weak. As long as ε is finite, we can make it as small as we please by weakening the coupling of the qubits to the bath, i.e., by rescaling the perturbation H_{SB} or equivalently by rescaling the field $\phi_\alpha(x, t)$.

What is the criterion for infrared convergence? Let us suppose that the qubits are uniformly distributed in D -dimensional space and that the bath fluctuations are “critical,” i.e., algebraically decaying in space and in time. We say that the scale dimension of the field ϕ is δ and the dynamical critical exponent is z if, for large-scale factor λ , the bath two-point function scales according to

$$\Delta(\lambda x_1, \lambda^z t_1; \lambda x_2, \lambda^z t_2) \sim \lambda^{-2\delta} \Delta(x_1, t_1; x_2, t_2), \quad (34)$$

thus the time t scales like z powers of the spatial distance x . This means that the integral of the two-point function scales as

$$\int dt d^D x |\Delta(x, t; 0, 0)| \sim R^{D+z-2\delta}, \quad (35)$$

where R is an infrared cutoff. Convergence in the infrared (finiteness of the limit $R \rightarrow \infty$) is ensured provided that

$$D + z < 2\delta. \quad (36)$$

If this criterion is satisfied, then scalable quantum computing is achievable at weak coupling. If it is not satisfied, then

scalable quantum computing might still be possible, but our version of the threshold theorem does not guarantee it. The same criterion was previously stated by Novais *et al.* [23,24] though without rigorous justification.

B. Almost-Markovian noise

The noise is Markovian if the bath immediately “forgets” any quantum information it receives, so that the information never returns to the system. Though this is never strictly the case, it can be true to an excellent approximation if the characteristic correlation time of the bath is very short compared to the time resolution with which we monitor the system’s behavior. In the Gaussian noise model, the noise is Markovian if the bath’s two-point correlation function is proportional to a delta function of the time difference,

$$\Delta(t_1, x_1; t_2, x_2) \propto \delta(t_1 - t_2). \quad (37)$$

We could say that the noise is “almost Markovian” if the correlation function Δ is a sharply peaked function of the time difference, e.g., with width τ_c much less than the duration t_0 of a single quantum gate. In that case, our expression for the noise strength becomes

$$\varepsilon^2 = C \max_{\text{loc}} \left(\int_{1,\text{loc}} \int_{2,\text{all}} |\bar{\Delta}(1,2)| \right) \approx \Gamma t_0. \quad (38)$$

For each fixed value of t_1 , the sharply peaked t_2 integral generates the factor Γ , and then integrating t_1 over the duration of the location generates the factor t_0 .

We may interpret Γ as an error rate per unit time and Γt_0 as an error probability per gate. But note that the noise strength ε is not this error rate, but rather its square root $(\Gamma t_0)^{1/2}$, in effect the amplitude of the error. In the Markovian case, fault paths really *do* decohere, and errors *can* be assigned probabilities rather than amplitudes. But our derivation of Eq. (27) is too general and insufficiently clever to exploit this property; hence our threshold condition requires the error amplitude rather than its square to be less than ε_0 .

Despite this deficiency, at least our threshold criterion for local noise, when applied to the almost-Markovian case, improves on the operator norm criterion Eq. (11). If the two-point function has a narrow peak of width τ_c whose integral is Γ , then the *height* of the peak is of order Γ/τ_c , and this peak height can be interpreted as the norm squared of the system-bath Hamiltonian, as in Eq. (15). Thus Eq. (11) estimates the noise strength as

$$\varepsilon \sim \sqrt{\Gamma/\tau_c} t_0, \quad (39)$$

which is even more pessimistic than Eq. (38). The estimate Eq. (39) diverges as the ultraviolet cutoff τ_c^{-1} is removed. But the estimate Eq. (38) depends on the area under the peak rather than its height and so has a smooth limit as $\tau_c \rightarrow 0$.

C. Ohmic noise

To further explore the sensitivity to high-frequency noise of our estimated noise strength, let us consider the Ohmic case, as in Sec. II C. If the Fourier transform $\tilde{\Delta}(\omega)$ of the

two-point correlation function is given by Eq. (17), then the real-time correlation function is

$$\Delta(t) = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-i\omega t} \tilde{\Delta}(\omega) = \frac{-A}{(t - i\tau_c)^2}. \quad (40)$$

The function $\Delta(t)$ has a short-time singularity at $t=0$ that is regulated by the cutoff τ_c , but the real and imaginary parts of $\Delta(t)$ both oscillate, so that its time integral vanishes,

$$\int_{-\infty}^{\infty} dt \Delta(t) = \tilde{\Delta}(\omega=0) = 0. \quad (41)$$

However the estimated noise strength ε , which is required to be small by the threshold condition, depends on the integral of the *modulus* of $\Delta(t)$,

$$|\Delta(t)| = \frac{A}{t^2 + \tau_c^2}, \quad (42)$$

which is of course non-negative and has a nonvanishing time integral; the estimated noise strength is

$$\varepsilon = \left(C \int_{\text{loc}} dt \int_{\text{all}} ds |\Delta(t-s)| \right)^{1/2} = \sqrt{\pi C A} \left(\frac{t_0}{\tau_c} \right)^{1/2}. \quad (43)$$

This estimate is ultraviolet divergent, but comparing to Eq. (18) we see that the divergence has been improved from linear to square-root dependence on the ultraviolet cutoff τ_c^{-1} .

Despite the improvement, the surviving ultraviolet sensitivity in this estimate of ε (for the case of Ohmic noise) is troubling, as it significantly reduces the class of noise models for which we can conclude that quantum computing is scalable. Therefore it is important to understand the origin of the ultraviolet divergence. One might suspect at first that the ultraviolet sensitivity arises because the range of the dt integral in Eq. (43) is a window of width t_0 with a sharp boundary. But, in fact, for Ohmic noise the sharp boundary generates only a mild logarithmic ultraviolet divergence, not a power divergence. The actual reason for the power divergence is that we have pled complete ignorance regarding the frequency spectrum of the ideal system Hamiltonian H_S . Therefore, we are required to be maximally pessimistic about how the oscillating phase of the wave function arising from the ideal system dynamics matches with the phase of the bath fluctuations. For that reason our estimate of ε involves an integral of the modulus of $\Delta(t)$ rather than $\Delta(t)$ itself.

With further assumptions about the ideal system dynamics we ought to be able to exclude this highly pessimistic scenario, leading to an estimate of ε with milder ultraviolet sensitivity. A natural idea is to attempt a “renormalization-group improvement” of the noise model, that is, to “coarse grain” in time, stretching the short-time cutoff τ_c , while adjusting the bath fluctuations to keep invariant the effect of the noise on the system. Formally Ohmic noise is “marginal,” meaning that the naive renormalization-group scaling generates only logarithmic cutoff dependence, not the square-root dependence found in Eq. (43). However, rigorously justifying this naive scaling turns out to be technically difficult in part because H_{SB} couples the system operators to *unbounded* bath operators in the Gaussian noise model. It might be in-

interesting to see if further technical assumptions (which one would hope to justify *a posteriori*) about the system-bath state $|\Psi_{SB}(t)\rangle$ during the course of the computation would lead to a less ultraviolet-sensitive threshold condition, but we have not yet succeeded in finding useful results with this character.

V. DERIVATION

In this section we will derive Eq. (27). Our task is to estimate a value of ε such that

$$\begin{aligned} \|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r)\|^2 &= \langle \Psi_{SB}^{\text{bad}}(\mathcal{I}_r) | \Psi_{SB}^{\text{bad}}(\mathcal{I}_r) \rangle \\ &= \langle \Psi_{SB}^0 | U_{SB}^{\text{bad}}(\mathcal{I}_r)^\dagger U_{SB}^{\text{bad}}(\mathcal{I}_r) | \Psi_{SB}^0 \rangle \leq \varepsilon^{2r} \end{aligned} \quad (44)$$

[see Eq. (9)]. Here U_{SB} is the joint system-bath time-evolution operator from the beginning of the computation until just before the measurements that will read out the final result, and $U_{SB}^{\text{bad}}(\mathcal{I}_r)$ denotes the sum of all the terms in the perturbation expansion of U_{SB} such that the perturbation H_{SB} is inserted at least once in each of the r specified locations in the set \mathcal{I}_r . The initial state of the system and bath is assumed to be the pure product state $|\Psi_{SB}^0\rangle = |\Psi_S^0\rangle \otimes |\Psi_B^0\rangle$, where the bath's state $|\Psi_B^0\rangle$ is Gaussian; that is, the expectation values of the bath operators $\{\phi_\alpha(x, t)\}$ obey Gaussian statistics in this state. For now we assume that system qubits are prepared only at the start of the computation and measured only at the end, with evolution governed by the Hamiltonian $H = H_S + H_B + H_{SB}$ in between; this assumption can be relaxed, as we will discuss in Sec. VI.

A. Keldysh diagrams

The terms in the perturbation expansion can be associated with diagrams, where in each diagram the perturbation H_{SB} is inserted at a specified set of points in spacetime. We may think of the sum of these diagrams as representing the expectation in the state $|\Psi_{SB}^0\rangle$ of the product of the forward evolution operator of the system and bath (i.e., U_{SB}), from the initial to the final time, followed by the backward evolution operator (i.e., U_{SB}^\dagger), from the final to the initial time. It is convenient to fold the diagram into a hairpin shape as in Fig. 2, so that the diagram has two branches that are aligned in time. The upper branch represents the evolution forward in time; here the inserted perturbations are “time-ordered,” meaning that operators inserted at later times act after operators inserted at earlier times. The lower branch represents the evolution backward in time; here the inserted perturbations are “anti-time-ordered,” meaning that operators inserted at earlier times act after operators inserted at later times. Furthermore, all operators inserted on the lower branch act after operators inserted on the upper branch. Diagrams with this structure are sometimes called “Keldysh diagrams.”

In each diagram, the evolution of system and bath is governed by the uncoupled Hamiltonian $H_0 = H_S + H_B$ in between successive insertions of H_{SB} . Since $|\Psi_{SB}^0\rangle$ is a product state, the diagram's contribution to the expectation value factorizes into the product of a system expectation value and a bath expectation value. Consider a diagram where the operator $\sigma_{\alpha_j} \otimes \tilde{\phi}_{\beta_j}$ is inserted on the upper branch acting on qubit x_j at

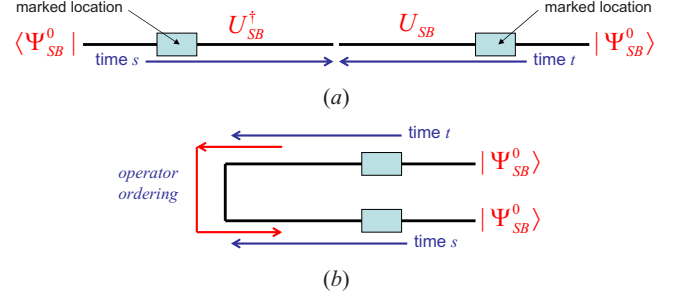


FIG. 2. (Color online) The diagram in (a) represents the norm squared of the system-bath state $U_{SB}|\Psi_{SB}^0\rangle$. By bending it into a hairpin shape we obtain the “Keldysh diagram” shown in (b). A “marked location” in the circuit appears twice in the Keldysh diagram, once on the upper branch (as a contribution to U_{SB}) and once on the lower branch (as a contribution to U_{SB}^\dagger).

time t_j , for $j = 1, 2, 3, \dots, n$, and the operator $\sigma_{\beta_k} \otimes \tilde{\phi}_{\beta_k}$ is inserted on the lower branch acting on qubit y_k at time s_k , for $k = 1, 2, 3, \dots, m$. Taking into account the uncoupled evolution in between insertions and the Keldysh operator ordering rules (where $t_n > t_{n-1} > \dots > t_1$ and $s_m < s_{m-1} < \dots < s_1$), this diagram's contribution to the expectation value $\langle \Psi_{SB}^0 | U_{SB}^\dagger U_{SB} | \Psi_{SB}^0 \rangle$ is

$$\begin{aligned} & i^m (-i)^n \langle \Psi_S^0 | \sigma_{\beta_m}(y_m, s_m) \cdots \sigma_{\beta_1}(y_1, s_1) \\ & \quad \times \sigma_{\alpha_n}(x_n, t_n) \cdots \sigma_{\alpha_1}(x_1, t_1) | \Psi_S^0 \rangle \\ & \quad \times \langle \Psi_B^0 | \phi_{\beta_m}(y_m, s_m) \cdots \phi_{\beta_1}(y_1, s_1) \\ & \quad \times \phi_{\alpha_n}(x_n, t_n) \cdots \phi_{\alpha_1}(x_1, t_1) | \Psi_B^0 \rangle. \end{aligned} \quad (45)$$

Here $\sigma_\alpha(x, t) = U_S(t)^\dagger \sigma_\alpha(x) U_S(t)$ and $\phi_\alpha(x, t) = U_B(t)^\dagger \tilde{\phi}_\alpha(x, t) U_B(t)$ are the “interaction-picture” operators that evolve according to the uncoupled system-bath dynamics. Using the Gaussian statistics (i.e., “Wick's theorem”), the bath expectation value can be expressed as a sum of products of Keldysh-ordered two-point correlation functions. Summing $\{\alpha_1, \alpha_2, \dots, \alpha_n\}$ and $\{\beta_1, \beta_2, \dots, \beta_m\}$ from 1 to 3, summing $\{x_1, x_2, \dots, x_n\}$ and $\{y_1, y_2, \dots, y_m\}$ over all qubits, integrating $\{t_1, t_2, \dots, t_n\}$ and $\{s_1, s_2, \dots, s_m\}$ over the interval from the initial to the final time, and finally summing n and m from 0 to ∞ , we would recover the full expectation value $\langle \Psi_{SB}^0 | U_{SB}^\dagger U_{SB} | \Psi_{SB}^0 \rangle = 1$. More precisely, to generate the full system-bath evolution operator U_{SB} , for each fixed n we sum $\{(x_1, t_1), (x_2, t_2), (x_3, t_3), \dots, (x_n, t_n)\}$ over all time-ordered sets of n space-time positions inside the circuit. This is equivalent to integrating each (x_j, t_j) over all spacetime, and then dividing by $n!$ to compensate for the overcounting of the sets (each set has been included $n!$ times). Similarly, to generate U_{SB}^\dagger , for each fixed m we sum $\{(y_1, s_1), (y_2, s_2), (y_3, s_3), \dots, (y_m, s_m)\}$ over all anti-time-ordered sets of m space-time positions inside the circuit.

B. One marked location

But we do not want to sum all the diagrams; instead the sum should include only those diagrams such that all of the r locations in the set \mathcal{I}_r are bad on both the upper and lower




FIG. 3. (Color online) A Keldysh diagram contributing to $\langle \Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=1}) | \Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=1}) \rangle$, where black dots are insertions, lines connecting pairs of dots are contractions, and the shaded rectangle on each branch indicates the one marked location. In this diagram, the marked location is bad because there is an insertion inside the marked location on each branch.

branches. Let us first consider the case $r=1$, where one particular circuit location has been “marked” as bad. To get a useful bound, it is helpful to organize this sum in a particular way. Because the marked location is bad, there must be at least one insertion of the perturbation inside this location on both the upper and lower branches, as in Fig. 3. Therefore, there must be an earliest insertion inside the marked location on each branch. Also, if the marked location is a two-qubit gate, then the earliest insertion could act on either one of the two qubits. For now, let us fix on each branch the time of the earliest insertion inside the marked location, the qubit on which the earliest insertion acts, and the corresponding Pauli operator. Later on we will integrate the time of the earliest insertion over the marked location and also sum over Pauli operators and the qubits at the location, but not yet.

With the earliest insertion fixed on each branch and after expanding the bath expectation value in terms of bath two-point functions, we can identify two classes of diagrams. In “class 1” diagrams, the earliest insertions on the two branches are contracted with one another and in “class 2” diagrams they are not, as shown in Fig. 4. We will find upper bounds on the sum of all the diagrams in each class.

1. Class 1 diagrams

First consider the class 1 diagrams. Each diagram in the class has as a factor the two-point function $\Delta(\beta, y, s; \alpha, x, t)$, where σ_α acts on qubit x at time t on the upper branch and




FIG. 4. (Color online) Skeleton Keldysh diagrams contributing to $\langle \Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=1}) | \Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=1}) \rangle$, showing the earliest insertion of H_{SB} inside the marked location on both branches. For the class 1 diagram shown in (a), the earliest insertions on the two branches are contracted with one another and for the class 2 diagram shown in (b), the earliest insertions are contracted with other insertions elsewhere. Other diagrams in each class are obtained by dressing the skeleton diagrams with additional insertions and contractions, except that no insertions are allowed inside the marked locations at times before the earliest insertion.

σ_β acts on qubit y at time s on the lower branch. The simplest diagram in the class, which we will call the “skeleton” diagram, has only two insertions and one contraction; its value is

$$\langle \Psi_S^0 | \sigma_\beta(y, s) \sigma_\alpha(x, t) | \Psi_S^0 \rangle \langle \Psi_B^0 | \phi_\beta(y, s) \phi_\alpha(x, t) | \Psi_B^0 \rangle. \quad (46)$$

The other class 1 diagrams are obtained by “dressing” this skeleton in all possible ways by adding further insertions and contractions.

However, remember that we have fixed t and s to be the times of the earliest insertions of the perturbation on the upper and lower branches, respectively. Therefore, an additional insertion is “legal” only if it avoids times earlier than t inside the marked location on the upper branch and times earlier than s inside the marked location on the lower branch. With this proviso, all class 1 diagrams arise when we dress the skeleton class 1 diagram with all possible additional legal insertions and contractions.

The class 1 diagrams can be summed up and expressed in a compact form. For this purpose we introduce what we call the “hybrid picture,” which is in a sense intermediate between the interaction and Heisenberg pictures. Let us define the “hybrid Hamiltonian” H^{hyb} by

$$H^{\text{hyb}} = \begin{cases} H_S + H_B & \text{in each marked location prior to the earliest insertion,} \\ H_S + H_B + H_{SB} & \text{elsewhere.} \end{cases} \quad (47)$$

That is, in the hybrid Hamiltonian, the perturbation H_{SB} “turns off” inside the marked location before time t on the upper branch and before time s on the lower branch. When we sum up all the legal insertions and contractions, the value [Eq. (46)] of the skeleton diagram is transformed into

$$\langle \Psi_{SB}^0 | \sigma_\beta^{\text{hyb}}(y, s) \sigma_\alpha^{\text{hyb}}(x, t) | \Psi_{SB}^0 \rangle \langle \Psi_B^0 | \phi_\beta(y, s) \phi_\alpha(x, t) | \Psi_B^0 \rangle. \quad (48)$$

Here the interaction-picture operator $\sigma_\alpha(x, t)$ has been replaced by the hybrid-picture operator $\sigma_\alpha^{\text{hyb}}(x, t)$

$= U_{SB}^{\text{hyb}}(t)^\dagger \sigma_\alpha(x) U_{SB}^{\text{hyb}}(t)$, and furthermore the expectation value of the system operator is now evaluated in the system-bath state $|\Psi_{SB}^0\rangle$ rather than the system state $|\Psi_S^0\rangle$. If the expression [Eq. (48)] is expanded in powers of H_{SB}^{hyb} , all the legal insertions and only the legal insertions are generated. For each choice of insertions, evaluating the bath expectation value using Wick’s theorem yields a sum over all the contractions included in class 1. Thus, the assumption that the bath fluctuations are Gaussian is crucial for the derivation of Eq. (48).

Therefore, we obtain an exact expression for the sum of all the diagrams in class 1 from Eq. (48) by now integrating the earliest insertions on both branches over the marked location, finding

$$\begin{aligned} & \sum \text{class 1 diagrams} \\ &= \sum_{x,y \in \text{loc}} \int_{\text{loc}} dt \int_{\text{loc}} ds \sum_{\alpha,\beta} \langle \Psi_{SB}^0 | \sigma_{\beta}^{\text{hyb}}(y,s) \sigma_{\alpha}^{\text{hyb}}(x,t) | \Psi_{SB}^0 \rangle \\ & \quad \times \langle \Psi_B^0 | \phi_{\beta}(y,s) \phi_{\alpha}(x,t) | \Psi_B^0 \rangle. \end{aligned} \quad (49)$$

Now, the operator $\sigma_{\alpha}^{\text{hyb}}(x,t)$ differs from the Pauli operator $\sigma_{\alpha}(x)$ by a mere unitary change of basis and therefore has sup norm $\|\sigma_{\alpha}^{\text{hyb}}(x,t)\|=1$. From Eq. (49) we then conclude that

$$\begin{aligned} & \left| \sum \text{class 1 diagrams} \right| \\ & \leq \sum_{x,y \in \text{loc}} \int_{\text{loc}} dt \int_{\text{loc}} ds \sum_{\alpha,\beta} |\langle \Psi_B^0 | \phi_{\beta}(y,s) \phi_{\alpha}(x,t) | \Psi_B^0 \rangle| \\ & = \int_{1,\text{loc}} \int_{2,\text{loc}} |\bar{\Delta}(1,2)| \end{aligned} \quad (50)$$

in the notation of Eq. (24). This is our bound on the sum of all class 1 diagrams.

Note that in Eq. (49) the integrand is the product of a bath two-point correlation function and a ‘‘hybridized’’ system two-point correlation function. If the bath correlation function has a high-frequency component and the system correlation function does not, then the contribution to the time integral arising from the high-frequency bath fluctuations may be strongly suppressed. But the estimate in Eq. (50) is very crude—it applies irrespective of the frequency spectrum of the system correlation function—and we could get a better estimate if we assumed that the system correlation function has little power at high frequency. Furthermore, such an assumption seems physically reasonable; the natural frequencies of the system dynamics are set by the energy splitting of the logical states and by the characteristic time scale (e.g., the gate duration t_0) on which the time-dependent system Hamiltonian varies. Unfortunately, though, finding a rigorous bound on the high-frequency hybrid system correlation function is not trivial because the hybrid Hamiltonian includes the system-bath coupling H_{SB} , an unbounded operator. If the bath has a low temperature, then we expect that high-frequency bath oscillators are likely to be in their ground states, but to prove a threshold theorem, we need to rule out relatively unlikely events that might foil the computation. That is not so easy to do, especially if the Hamiltonian is unbounded. So in this paper we will mostly pursue the consequences of the crude estimate Eq. (50) and other similar estimates, leaving for future work the challenge of improving the results via tighter bounds on the integral in Eq. (49). However, we *can* obtain a stronger bound for the case of pure dephasing noise, discussed in Sec. VII.

To prevent confusion, we remark that our hybrid picture is a rather strange concept, in that the Hamiltonian that governs evolution on the upper branch of the Keldysh diagram is

different than the Hamiltonian for the lower branch. If we were using Keldysh diagrams the way they are usually used, to track the evolution of the system’s density operator, this feature would be unacceptable because time evolution would not preserve the density operator’s trace. For us, though, the hybrid Hamiltonian is merely a technical trick for bounding the sum of a class of diagrams and should not be interpreted literally as the Hamiltonian of a physical system.

2. Class 2 diagrams

Now consider the class 2 diagrams. The earliest insertions inside the marked location on the upper and lower branches of the Keldysh diagram are not contracted with one another; rather each is contracted with an insertion at another location. Let us say that the earliest insertion at (x,t) in the marked location on the upper branch is contracted with an insertion at space-time position (z,u) , which could be on either the upper or lower branch, and that the earliest insertion at (y,s) in the marked location on the lower branch is contracted with an insertion at (w,v) , which also could be on either the upper or lower branch. In principle z, w could be the spatial labels of any two qubits in the computer and u, v could be any time between the initial and final times, *except* that the insertions at (z,u) and (w,v) must be legal, that is, neither can be inside the marked location on the upper branch earlier than t or inside the marked location on the lower branch earlier than s .

For the class 2 diagrams, let us for now imagine fixing the insertions at (z,u) and at (w,v) that are contracted with the earliest insertions; we will integrate over these spacetime positions later on. The simplest diagram in the class, the skeleton diagram with only two contractions, has the value [except for a phase factor that depends on the choice of branch for the insertions at (z,u) and (w,v)]

$$\begin{aligned} & \langle \Psi_S^0 | T^*(\sigma_{\beta}(y,s) \sigma_{\delta}(w,v) \sigma_{\gamma}(z,u) \sigma_{\alpha}(x,t)) | \Psi_S^0 \rangle \\ & \quad \times \langle \Psi_B^0 | T^*(\phi_{\beta}(y,s) \phi_{\delta}(w,v)) | \Psi_B^0 \rangle \\ & \quad \times \langle \Psi_B^0 | T^*(\phi_{\gamma}(z,u) \phi_{\alpha}(x,t)) | \Psi_B^0 \rangle, \end{aligned} \quad (51)$$

where T^* denotes the proper Keldysh ordering. Other diagrams in class 2 are obtained by dressing this skeleton with additional insertions and contractions in all possible legal ways. As in our discussion of the class 1 diagrams, summing all the ways to dress the skeleton transforms the interaction-picture system operators into hybrid-picture operators, yielding (up to a phase)

$$\begin{aligned} & \langle \Psi_{SB}^0 | T^*(\sigma_{\beta}^{\text{hyb}}(y,s) \sigma_{\delta}^{\text{hyb}}(w,v) \sigma_{\gamma}^{\text{hyb}}(z,u) \sigma_{\alpha}^{\text{hyb}}(x,t)) | \Psi_{SB}^0 \rangle \\ & \quad \times \langle \Psi_B^0 | T^*(\phi_{\beta}(y,s) \phi_{\delta}(w,v)) | \Psi_B^0 \rangle \\ & \quad \times \langle \Psi_B^0 | T^*(\phi_{\gamma}(z,u) \phi_{\alpha}(x,t)) | \Psi_B^0 \rangle. \end{aligned} \quad (52)$$

To obtain the sum of all class 2 diagrams, we now sum over Pauli operator labels and spacetime positions, obtaining

\sum class 2 diagrams

$$\begin{aligned}
 &= \sum_{x,y \in \text{loc}} \sum_{z,w \in \text{all}'} \int_{\text{loc}} dt \int_{\text{loc}} ds \int_{\text{all}'} du \int_{\text{all}'} dv \sum_{\alpha,\beta,\gamma,\delta} (\text{phase}) \langle \Psi_{SB}^0 | T^*(\sigma_{\beta}^{\text{hyb}}(y,s) \sigma_{\delta}^{\text{hyb}}(w,v) \sigma_{\gamma}^{\text{hyb}}(z,u) \sigma_{\alpha}^{\text{hyb}}(x,t)) | \Psi_{SB}^0 \rangle \\
 &\quad \times \langle \Psi_B^0 | T^*(\phi_{\beta}(y,s) \phi_{\delta}(w,v)) | \Psi_B^0 \rangle \langle \Psi_B^0 | T^*(\phi_{\gamma}(z,u) \phi_{\alpha}(x,t)) | \Psi_B^0 \rangle.
 \end{aligned} \tag{53}$$

Here the notation all' indicates that the qubit positions z and w are summed over *both* branches of the Keldysh diagram and that the times u and v are also integrated over both branches. Furthermore, it is understood that the integral over u and v is restricted to legal insertions (times in the upper-branch marked location earlier than t and in the lower-branch marked location earlier than s are excluded).

As for the class 1 diagrams, we obtain a bound on the sum of class 2 diagrams by noting that the expectation value of the product of system operators has modulus not larger than 1, finding

$$\left| \sum \text{class 2 diagrams} \right| \leq \left(2 \int_{1,\text{loc}} \int_{2,\text{all}} |\bar{\Delta}(1,2)| \right)^2. \tag{54}$$

To obtain Eq. (54) we have noted that the Keldysh ordering is irrelevant when we take the modulus of the bath two-point function, and that in the sum of the moduli of all diagrams we can extend the integral over legal insertions to an integral over all insertions to obtain an upper bound. Here the notation all indicates that the second leg of the correlation function is summed over all qubits and integrated over all times; the factor of 2 accompanies the integral $\int_{2,\text{all}}$ because the insertions at (z,u) and at (w,v) can be on either one of the two branches of the Keldysh diagram.

For our upper bound on the sum of class 1 diagrams, both legs of the bath's two-point function are integrated over the marked location, while in the upper bound on the sum of class 2 diagrams, one leg is integrated over the marked location, while the other is integrated over all qubits and all times. This distinction is not so important if the spatial and temporal correlations decay rapidly, but it can be quite important if the decay is slow, as we have already discussed in Sec. IV A. The upper bound on the sum of class 1 diagrams is still valid, though weaker, if we extend the integral for one of the legs from the marked location to all of spacetime. Then by adding together the contributions from diagrams of both classes, we find

$$\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=1})\|^2 \leq E + 4E^2, \tag{55}$$

where

$$E = \max_{\text{loc}} \left(\int_{1,\text{loc}} \int_{2,\text{all}} |\bar{\Delta}(1,2)| \right). \tag{56}$$

If E is small (the typical case of interest), then the class 1 diagrams dominate and the contribution from class 2 diagrams is higher order in E . We emphasize again that the integral $\int_{2,\text{all}}$ in the definition of E is confined to a single

branch of the Keldysh diagram and that the factor of 2 in Eq. (54) arises because the insertion inside the marked location can be contracted with an insertion on either branch.

C. Many marked locations

Now we want to consider the case where there are r marked locations. The perturbation H_{SB} must be inserted at least once in each of the r marked locations on both the upper and lower branches of the Keldysh diagram. In order to get an upper bound on the sum of all such diagrams, we will organize the sum following the same ideas as in our discussion of the $r=1$ case. In each marked location on each branch, there must be an earliest insertion of the perturbation, and this earliest insertion is contracted with another insertion elsewhere, which could be on either branch.

A skeleton graph contains a “minimal” set of contractions—each contraction in the skeleton has at least one leg attached to the earliest insertion in a marked location. We distinguish two types of contractions in the skeleton: an “internal” contraction links two earliest insertions and an “external” contraction links an earliest insertion with another legal insertion which is not an earliest insertion. The skeleton diagrams can be classified according to the number k of internal contractions. If there are r marked locations, and therefore all together $2r$ marked locations between the two branches, then k can vary from 0 to r ; if there are k internal contractions then there are $2(r-k)$ external contractions. For $r=1$, a skeleton diagram in what we called class 1 has $k=1$ internal contractions, and a skeleton diagram in class 2 has $k=0$ internal contractions. For $r=2$, the ten distinct skeleton diagrams are shown in Fig. 5. There are three diagrams with two internal contractions, six diagrams with one internal contraction, and one diagram with no internal contractions.

The value of a skeleton diagram, with all insertions and contractions fixed, can be expressed as a product of the ex-




FIG. 5. (Color online) Skeleton Keldysh diagrams contributing to $\langle \Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=2}) | \Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=2}) \rangle$, showing the earliest insertion of H_{SB} inside each of two marked locations on both branches. There are three skeleton diagrams with two internal contractions, six skeleton diagrams with one internal contraction, and one skeleton diagram with no internal contraction, where we say that a contraction is internal if it links two earliest insertions.

pectation value of a string of Keldysh-ordered interaction-picture system operators

$$\langle \Psi_S^0 | T^* [\sigma(i_1) \cdots \sigma(i_k) \sigma(j_1) \cdots \sigma(j_k) \sigma(m_1) \cdots \sigma(m_{2(r-k)}) \sigma(n_1) \cdots \sigma(n_{2(r-k)})] | \Psi_S^0 \rangle \quad (57)$$

and a product of Keldysh-ordered bath two-point functions

$$\begin{aligned} & (\text{phase}) \times \langle \Psi_B^0 | T^* [\phi(i_1) \phi(j_1)] | \Psi_B^0 \rangle \\ & \cdots \langle \Psi_B^0 | T^* [\phi(i_k) \phi(j_k)] | \Psi_B^0 \rangle \\ & \times \langle \Psi_B^0 | T^* [\phi(m_1) \phi(n_1)] | \Psi_B^0 \rangle \\ & \cdots \langle \Psi_B^0 | T^* [\phi(m_{2(r-k)}) \phi(n_{2(r-k)})] | \Psi_B^0 \rangle. \end{aligned} \quad (58)$$

Here we have attached labels $1, 2, 3, \dots, 2r$ to the $2r$ marked locations on the two branches, and, e.g., $\sigma(i_1)$ is shorthand for $\sigma_{\alpha_{i_1}}(x_{i_1}, t_{i_1})$, where (x_{i_1}, t_{i_1}) is the spacetime position of the first insertion inside marked location number i_1 . In Eqs. (57) and (58), locations i_1 through i_k are internally contracted with locations j_1 through j_k , while the remaining earliest insertions in locations m_1 to $m_{2(r-k)}$ are contracted with insertions labeled n_1 through $n_{2(r-k)}$ which are not earliest insertions.

When we sum up all ways to dress the skeleton with additional insertions and contractions, we obtain an expression of the same form, but with the interaction-picture system operators replaced by hybrid-picture system operators and $|\Psi_S^0\rangle$ replaced by $|\Psi_{SB}^0\rangle$. Bounding the system operator expectation value by one and summing over the Pauli operator labels, we obtain a bound on the sum of all dressed skeleton diagrams,

$$\left| \sum \text{dressed skeletons} \right| \leq \prod_{a=1}^k |\bar{\Delta}(i_a, j_a)| \prod_{b=1}^{2(r-k)} |\bar{\Delta}(m_b, n_b)|. \quad (59)$$

Now, keeping fixed the choice of which locations are internally contracted with one another, we can integrate each i_a , j_a , and m_b over the specified marked location, while integrating n_b over all locations on both branches. The integral is bounded above by

$$\int \left| \sum \text{dressed skeletons} \right| \leq \prod_{a=1}^k G(i_a, j_a) \prod_{b=1}^{2(r-k)} 2E(m_b), \quad (60)$$

where

$$\begin{aligned} G(i_a, j_a) &= \int_{1, \text{loc}(i_a)} \int_{2, \text{loc}(j_a)} |\bar{\Delta}(1, 2)|, \\ E(m_b) &= \int_{1, \text{loc}(m_b)} \int_{2, \text{all}} |\bar{\Delta}(1, 2)|, \end{aligned} \quad (61)$$

and where the factor of 2 multiplying $E(m_b)$ results from summing n_b over both branches.

Now, with the number k of internal contractions still fixed, we can sum Eq. (60) over all the ways that k contracted pairs of locations can be chosen from among $2r$ locations. We note that

$$\sum_{\text{contractions}(k)} \left(\prod_{a=1}^k G(i_a, j_a) \right) \leq \frac{1}{k!} \left(\sum_{\substack{i,j=1 \\ i < j}}^{2r} G(i, j) \right)^k. \quad (62)$$

This inequality holds because $[\sum_{i,j=1}^{2r} G(i, j)]^k$ contains the term corresponding to each contraction $k!$ times and also contains other non-negative terms. Furthermore,

$$\sum_{\substack{i,j=1 \\ i < j}}^{2r} G(i, j) \leq \sum_{i=1}^{2r} E(i) \quad (63)$$

because the expression on the right-hand side contains all the terms on the left-hand side, plus other non-negative terms. We conclude that

$$\sum_{\text{contractions}(k)} \int \left| \sum \text{dressed skeletons} \right| \leq \frac{1}{k!} (2rE)^k (2E)^{2(r-k)}, \quad (64)$$

with E defined as in Eq. (56).

It remains to sum over k the number of internal contractions,

$$\begin{aligned} \left| \sum \text{diagrams} \right| &\leq \sum_{k=0}^r \left(\sum_{\text{contractions}(k)} \int \left| \sum \text{dressed skeletons} \right| \right) \\ &\leq \sum_{k=0}^r \frac{r^k}{k!} (2E)^{2r-k}. \end{aligned} \quad (65)$$

Therefore, if we assume that $2E \leq 1$,

$$\begin{aligned} \|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r)\|^2 &\leq (2E)^r \sum_{k=0}^r \frac{r^k}{k!} (2E)^{r-k} \leq (2E)^r \sum_{k=0}^{\infty} \frac{r^k}{k!} \\ &= (2eE)^r = \varepsilon^{2r}, \end{aligned} \quad (66)$$

where

$$\varepsilon = \sqrt{2eE} \approx 2.34\sqrt{E}. \quad (67)$$

Thus we have derived Eq. (27). We note that Eq. (66) also applies for $r=1$, and in that case is weaker than the upper bound $\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_{r=1})\|^2 \leq E + 4E^2 = E(1 + 4E) \leq 3E$ found in Eq. (55), assuming $2E \leq 1$. We have made the assumption $2E \leq 1$ merely to simplify the derivation of the upper bound in Eq. (66). This assumption is harmless because the threshold value of E is much less than $1/2$.

VI. GENERALIZATIONS

A. Initial state of the bath

In our analysis, we have found it convenient to assume that the initial state of the bath is a pure state, but the analysis also applies if the bath starts out in a mixed state. Actu-

ally, we can “purify” a mixed state ρ_B^0 of the bath by introducing a fictitious reference system R and choosing the pure state $|\Psi_{BR}^0\rangle$ of BR so that

$$\rho_B^0 = \text{tr}_R(|\Psi_{BR}^0\rangle\langle\Psi_{BR}^0|). \quad (68)$$

Our previous analysis then applies, if we consider BR to be an “extended” bath, such that the system interacts with only the subsystem B of the extended bath.

However, the state $|\Psi_{BR}^0\rangle$ is not arbitrary; for our argument to apply it must be chosen so that the interaction-picture free field $\phi_\alpha(x,t)$ has Gaussian statistics in this state. For this it suffices for $|\Psi_{BR}^0\rangle$ to be an *undisplaced Gaussian squeezed state*. If we consider the reference system R , like the bath B , to be a system of uncoupled oscillators, then an undisplaced Gaussian squeezed state is obtained by applying a unitary transformation V to the oscillator ground state $|0_B, 0_R\rangle$, where the action of V on the annihilation operators is homogeneous and linear,

$$V^{-1}a_kV = \sum_j M_{kj}a_j + \sum_j N_{kj}a_j^\dagger. \quad (69)$$

Here the set $\{a_k\}$ includes annihilation operators for both the B oscillators and the R oscillators, and the matrices M and N obey constraints that ensure preservation of the commutation relations. V satisfies Eq. (69) if its logarithm is strictly quadratic in creation and annihilation operators, with no linear term. A special case is the thermal state of the bath, whose purification can be written as

$$\begin{aligned} & \exp\left(\sum_k r_k(a_{B,k}^\dagger a_{R,k}^\dagger - a_{B,k} a_{R,k})\right)|0_B, 0_R\rangle \\ &= \bigotimes_k \left(\sqrt{1 - \gamma_k^2} \sum_{n_k=0}^\infty \gamma_k^{n_k} |(n_k)_B, (n_k)_R\rangle \right), \end{aligned} \quad (70)$$

where $\gamma_k^2 = \tanh^2 r_k = e^{-2\beta\omega_k}$ and β is the inverse temperature. But our arguments apply to any Gaussian state $V|0_B, 0_R\rangle$ since the action of V in Eq. (69) maps free fields to new free fields that still satisfy Wick’s theorem and have mean zero.

B. Measurement and entropy removal

For fault-tolerant computing to work, there must be a mechanism for flushing the entropy introduced by noise. In the scheme we have analyzed, entropy is removed from the computer because error-correction gadgets use a supply of fresh ancilla qubits that are discarded after use. It has been understood that the initial state $|\Psi_S^0\rangle$ of the system includes all of the ancilla qubits that will be needed during the full course of the computation. But to model the actual situation, in which ancilla qubits are prepared as needed just before being used, we imagine that ancilla qubits are perfectly isolated from the bath until “opened” at the onset of the gadget in which they participate. Similarly, we imagine that the measurements of all ancilla qubits are delayed until the very end of the computation, but that these qubits are “closed” (their coupling to the bath is turned off) at the conclusion of the gadget in which they participate. With these stipulations, our noise model is equivalent to one in which ancilla qubits are repeatedly measured, reset, and reused.

We model the noisy preparation of an ancilla qubit as an ideal preparation followed by interaction with the bath for a specified duration. Since the state of the bath may evolve during the computation, the noise in the preparation may also depend on when the qubit is prepared. Still, we are taking it for granted that “pretty good” fresh ancillas can be prepared at any time or equivalently that qubits can be effectively erased at any time. Implicitly, we have adopted a “two-reservoir” hypothesis. One reservoir, which we have called the “bath,” interacts with the system qubits, causing noise. The other reservoir is the entropy “sink,” which carries away heat each time a qubit is erased. In our model, the bath and the sink are uncoupled, and the sink has infinite heat capacity—it never heats up no matter how many qubit erasures occur.

Because the bath interacts with the system, in principle, it might be driven far from its initial state in a manner that depends on the ideal computation being simulated. Our arguments have shown that, at least if the bath is a system of uncoupled oscillators and its initial state is Gaussian, the bath will not be pushed to a highly adversarial state that overpowers our efforts to make the computation robust. One wonders how that conclusion could be altered if we relax the two-reservoir hypothesis by coupling the sink and the bath or by eliminating the sink entirely. For example, we could attempt to model measurement and erasure more realistically by including entropy flow from the system to the bath. In that case, a bath of unbounded heat capacity would be needed to remove entropy from a noisy computation of unbounded size, and our modeling would need to incorporate a mechanism for equilibration of the bath. The goal would be to specify conditions under which the entropy flow from system to bath can be maintained well enough to support scalable quantum computation. For now, we put aside this ambitious project as an open problem for future consideration.

C. Postselection

Some fault-tolerant gadgets include *postselection*. For example, a gadget might consume a disposable piece of “quantum software,” an encoded ancilla state that is prepared offline and verified before coming into contact with the encoded data processed by the computation. The verification procedure includes measurements that check the accuracy of the preparation of the software, and the software is accepted only if the measurements have suitable outcomes; otherwise the software is rejected and the preparation is repeated. Therefore, estimates of the reliability of gadgets are conditioned on acceptance of the software, which is said to have a “post-selected” state.

Some fault-tolerant protocols, such as those analyzed in [10–12] make “extreme” use of postselection, meaning that the software is usually rejected and the preparation is typically repeated many times before it finally succeeds. For such protocols, noise with adversarial correlations can be a formidable foe since the adversary is empowered to enhance the probability of acceptance for atypical fault paths that are especially damaging. Thus, the threshold estimates based on extreme postselection proved in [10,11] apply for indepen-

dent noise but not for local stochastic noise. But other protocols, such as those analyzed in [6,8,25], make only “modest” use of postselection, meaning that software is accepted with reasonably high probability. For such protocols, a gadget’s failure rate, conditioned on acceptance of the software, can be easily estimated using the Bayes rule, even for the case of local stochastic noise.

The threshold estimate for local noise whose proof is sketched in the Appendix applies to a protocol with no postselection at all. For local noise, as for local stochastic noise, we do not know how to extend this proof to a protocol with extreme postselection. But it can be extended to a protocol with modest postselection. This observation is useful because threshold estimates based on protocols with modest postselection are typically higher than estimates based on protocols without postselection.

Before considering the case of local coherent noise, we recall how protocols with modest postselection can be analyzed for the case of local stochastic noise [6]; to be concrete, we will discuss the case where the level-1 gadgets are based on a quantum error-correcting code that corrects one error in a block. A properly designed gadget processes encoded data correctly if the software is accepted and the gadget contains no more than one fault. Therefore, the *joint* probability P_{joint} of acceptance of the software and failure of the gadget is bounded above by $B\varepsilon^2 + D\varepsilon^3$ for local stochastic noise with strength ε , where B is the number of malignant pairs of locations in the gadget where faults can cause failure (assuming the software is accepted) and D is the total number of sets of three locations in the gadget. On the other hand, the software will surely be accepted if there are no faults in the software preparation circuit, so the probability of acceptance P_{accept} is bounded below by $1 - C\varepsilon$, where C is the total number of locations in the circuit for software preparation and verification. Using the Bayes rule, we obtain an upper bound on the probability $P_{\text{conditional}}$ of failure conditioned on acceptance,

$$P_{\text{conditional}} = \frac{P_{\text{joint}}}{P_{\text{accept}}} \leq \frac{B\varepsilon^2 + D\varepsilon^3}{1 - C\varepsilon} \leq \varepsilon^2/\varepsilon_0 = \varepsilon^{(1)}, \quad (71)$$

where

$$\varepsilon_0^{-1} = \frac{1}{2}(B + C)(1 + \sqrt{1 + 4D/(B + C)^2}) \quad (72)$$

is determined by solving the equation $(B\varepsilon_0^2 + D\varepsilon_0^3)/(1 - C\varepsilon_0) = \varepsilon_0$. This argument gives a useful result if our lower bound on P_{accept} is not too small. In practice, it is often the case that $C \ll B$ and therefore $C\varepsilon_0 \ll 1$, so that the “postselection correction” arising from division by P_{accept} is a small effect.

There is another way to describe this estimate that is more readily generalized to the case of local coherent noise and also clarifies why the estimate applies to adversarial local stochastic noise. Imagine that n software preparation and verification attempts are executed in parallel, where we label the attempts by an index $i=1,2,3,\dots,n$ and suppose for the moment that the noise is uncorrelated. Now we distinguish $n+1$ possible ways for the gadget to be bad, depending on which preparation attempt (if any) is the first to be

accepted. If ancilla 1 is accepted, then the gadget fails with probability P_{joint} . But ancilla 1 is rejected with probability $P_{\text{reject}} = 1 - P_{\text{accept}}$, so the probability that ancilla 1 is rejected, ancilla 2 is accepted, and the gadget fails is $P_{\text{reject}}P_{\text{joint}}$. Similarly, the probability that ancilla m is the first to be accepted and the gadget fails is $P_{\text{reject}}^{m-1}P_{\text{joint}}$, and the probability that all n ancillas are rejected is P_{reject}^n . Summing the probability of all failure scenarios, we find

$$\begin{aligned} P_{\text{fail}} &= P_{\text{joint}} \left(\sum_{m=1}^n P_{\text{reject}}^{m-1} \right) + P_{\text{reject}}^n \\ &= \frac{P_{\text{joint}}}{P_{\text{accept}}} (1 - P_{\text{reject}}^n) + P_{\text{reject}}^n \\ &= \frac{P_{\text{joint}}}{P_{\text{accept}}} + P_{\text{reject}}^n \left(1 - \frac{P_{\text{joint}}}{P_{\text{accept}}} \right). \end{aligned} \quad (73)$$

In the limit $n \rightarrow \infty$, we recover the estimate Eq. (71), and even for $n=2$ we have $P_{\text{fail}} = O(\varepsilon^2)$. Furthermore, the upper bound on P_{fail} obtained from upper bounds on P_{joint} and P_{reject} applies not only to independent noise but also to correlated local stochastic noise—it can be regarded as an estimate of the effective noise strength $\varepsilon^{(1)}$ after one level-reduction step. For local stochastic noise, we sum over all failure scenarios at each of r marked locations and conclude that the probability that all r locations are bad is bounded above by $(\varepsilon^{(1)})^r$.

We can also apply this strategy of summing over all failure scenarios in the case of local coherent noise. First we note that, to preserve the framework assumed in Sec. V, we may imagine that all measurements in verification steps are postponed until the end of the computation. In the actual circuit, the “verification qubits” are measured inside gadgets, and then subsequent operations are conditioned on the classical measurement outcomes. To model this circuit in the framework where all measurements are postponed, we suppose that a verification qubit decouples from the bath at the time when it is measured in the actual circuit, and we replace operations conditioned on measurement outcomes by noiseless quantum gates conditioned on the state of the verification qubit, after decoupling from the bath but prior to being measured. Then we can estimate $\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(1)})\|^2$ by summing over $n+1$ failure scenarios at each of the r marked locations. In scenario 1, ancilla 1 is accepted and the gadget using ancilla 1 (including the preparation and verification of ancilla 1) has two or more faults. In scenario m , for $m=2,3,\dots,n$, the first $m-1$ ancillas are rejected, ancilla m is accepted, and the gadget using ancilla m has two or more faults. In scenario $n+1$, all n ancillas are rejected. Since the scenarios are perfectly distinguishable, they should be summed incoherently.

Now, in order for an ancilla to be rejected, there must be at least one fault in the circuit that prepares and verifies that ancilla. Therefore, in scenario m , we sum coherently over all fault paths such that there is at least one fault in each of the first $m-1$ ancilla preparation and verification circuits and at least two faults in the gadget using ancilla m . This sum includes all of the fault paths that contribute to the badness of the gadget under scenario m , but it also includes other fault paths that do not contribute to scenario m . However, since

the scenarios are distinguishable, there is no harm in including these additional unwanted scenarios if our goal is to obtain an upper bound on $\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(1)})\|^2$. This coherent sum for each scenario can be estimated by the method described in the Appendix. One finds that, for gadgets such that the post-selection correction to $\varepsilon^{(1)}$ is small in the case of local stochastic noise, the correction is small for local coherent noise as well.

In [25], the lower bound on the accuracy threshold $\varepsilon_0 \geq 1.94 \times 10^{-4}$ was established for local stochastic noise based on a protocol with modest use of postselection. Though we have not done the calculation in detail, we expect that a similar estimate $\varepsilon_0 \sim 10^{-4}$, based on the same protocol, also applies to the threshold noise strength for local coherent noise. (The argument in [9] achieves a higher threshold estimate for local stochastic noise but uses a different method that is less easily adapted to the case of local coherent noise.)

Of course, for the case of Gaussian noise, if gadgets include multiple parallel attempts to prepare and verify software, then all of these attempts should be included in the integral $\int_{2,\text{all}}$ in our estimate of the noise strength in Eq. (27).

D. Other considerations

It would be desirable to extend the derivation of our threshold result in several other directions. One possible approach is to allow the bath fluctuations to be weakly non-Gaussian by including small anharmonic corrections in the bath Hamiltonian H_B . But, though the effects of bath self-interactions can be analyzed perturbatively by standard methods, obtaining useful rigorous results summed to all orders of perturbation theory is not simple. Another worthy goal, already emphasized at the end of Sec. IV, is to formulate a threshold condition less sensitive to the high-frequency fluctuations of the bath, i.e., to noise with a frequency large compared to the natural frequencies of the ideal system Hamiltonian H_S . In principle this might be done by “integrating out” high-frequency noise, obtaining an effective noise model with a lower frequency cutoff that faithfully reproduces the impact of the noise on the simulated computation. Making such an analysis rigorous is another challenging open problem. In Sec. VII, though, we will discuss one special case in which an improved threshold estimate less sensitive to high-frequency noise can be achieved.

VII. DIAGONAL GAUSSIAN NOISE

As we discussed in Sec. V B 1, our general arguments do not place any constraints on the frequency spectrum of the hybrid-picture system operators. Therefore, we were forced to take the modulus of the bath two-point function in our estimate of the noise strength ε . As a result, our estimate has a sensitivity to high-frequency bath fluctuations that seems rather artificial.

There is at least one case where we have much better analytic control over the time dependence of the system operators, allowing us to obtain a better estimate of the noise strength that has milder sensitivity to high-frequency noise.

That is the case of pure dephasing noise, which we will discuss now.

In this noise model, the bath couples only to the z components of the qubits, so that the system-bath Hamiltonian is

$$H_{SB} = \sum_x \sigma_z(x) \otimes \tilde{\phi}(x,t), \tag{74}$$

where $\tilde{\phi}(x,t)$ is a Gaussian bath variable with mean zero. To further simplify the discussion (whose purpose is merely illustrative anyway), we will also assume there are no multi-qubit correlations in the noise (even though this might not be an accurate description of the noise in multiqubit gates). That is, we assume $\langle \phi(x,t)\phi(y,s) \rangle = 0$ for $x \neq y$, so that in effect each qubit is coupled to its own independent bath.

A scheme for fault-tolerant quantum computation customized for highly biased noise dominated by dephasing was formulated in [26] and further discussed in [27]. In this scheme, all gates are teleported. Furthermore the only fundamental operations used are single-qubit preparations, single-qubit measurements in the σ_x -eigenstate basis, and two-qubit controlled-phase (CPHASE) gates. A CPHASE gate is diagonal in the computational (i.e., σ_z -eigenstate) basis, with eigenvalues $(1,1,1,-1)$. Thus it can be realized by a time-dependent two-qubit system Hamiltonian that is also diagonal,

$$H_S = f(t)(\sigma_z \otimes \sigma_z - \sigma_z \otimes I - I \otimes \sigma_z), \tag{75}$$

where $\int dt f(t) = \pi/4$. This diagonal system Hamiltonian commutes with the system-bath Hamiltonian [Eq. (74)], whose action on the system qubits is diagonal. As in previous sections, we model a noisy qubit preparation as an ideal preparation followed by interaction with the oscillator bath, and we model a noisy measurement as interaction with the bath followed by an ideal measurement.

We can analyze the effect of the noise on the computation using interaction-picture perturbation theory, and in fact we can estimate a probability (rather than an amplitude) for the outcome of a qubit measurement to differ from the measurement outcome in the ideal computation. For each qubit, we distinguish between good diagrams, in which the perturbation H_{SB} is inserted an *even* number of times in between the (ideal) qubit preparation and the (ideal) qubit measurement, and the bad diagrams, in which the perturbation is inserted an *odd* number of times in between the preparation and the measurement. Because σ_z commutes with the ideal system Hamiltonian H_S and because $\sigma_z^2 = I$, in all good diagrams the outcome of the final σ_x measurement agrees with the result in the ideal quantum circuit, while in bad diagrams the measurement outcome is flipped.

Furthermore, the good and the bad parts of the system-bath state $|\Psi_{SB}\rangle$ are mutually orthogonal. To see this, imagine evaluating the inner product $\langle \Psi_{SB}^{\text{good}} | \Psi_{SB}^{\text{bad}} \rangle$ between the good and bad parts of the state for a single qubit and its associated bath. Since $|\Psi_{SB}^{\text{bad}}\rangle$ has an odd number of perturbation insertions and $|\Psi_{SB}^{\text{good}}\rangle$ has an even number, each Keldysh diagram contributing to $\langle \Psi_{SB}^{\text{good}} | \Psi_{SB}^{\text{bad}} \rangle$ is proportional to the expectation value of a product of an odd number




FIG. 6. (Color online) The three connected Keldysh diagrams for a single qubit subject to Gaussian dephasing noise. The first two diagrams are good because σ_z is inserted an even number of times on each branch and the third diagram is bad because σ_z is inserted an odd number of times on each branch.

of interaction-picture bath fields. All such diagrams vanish since $\phi(x, t)$ is Gaussian with mean zero. Because the good and bad parts of $|\Psi_{SB}\rangle$ are perfectly distinguishable, we can regard $\langle \Psi_{SB}^{\text{bad}} | \Psi_{SB}^{\text{bad}} \rangle$ as the probability of error in the final qubit measurement.

Let us compute this probability. The sum of all Keldysh diagrams (both good and bad) contributing to $\langle \Psi_{SB} | \Psi_{SB} \rangle$ for a single qubit is the exponential of the sum of “connected” diagrams. There are three connected diagrams, shown in Fig. 6. Thus

$$1 = \langle \Psi_{SB} | \Psi_{SB} \rangle = \exp(C_U + C_L + D). \quad (76)$$

Here,

$$C_U = - \int_{t>s} dt ds \langle \phi(t) \phi(s) \rangle \quad (77)$$

is the connected diagram in which two insertions on the upper Keldysh branch are contracted,

$$C_L = - \int_{t<s} dt ds \langle \phi(t) \phi(s) \rangle \quad (78)$$

is the connected diagram in which two insertions on the lower branch are contracted, and

$$D = \int dt ds \langle \phi(t) \phi(s) \rangle = -(C_U + C_L). \quad (79)$$

is the connected diagram in which an insertion on the upper branch is contracted with an insertion on the lower branch. (In all three diagrams, the factor due to the expectation value of the product of system operators is simply 1.) When Eq. (76) is expanded in powers of D , terms with an odd number of powers contribute to $\langle \Psi_{SB}^{\text{bad}} | \Psi_{SB}^{\text{bad}} \rangle$ because H_{SB} is inserted an odd number of times on each branch and terms with an even number of powers contribute to $\langle \Psi_{SB}^{\text{good}} | \Psi_{SB}^{\text{good}} \rangle$ because H_{SB} is inserted an even number of times on each branch. Thus,

$$\begin{aligned} P^{\text{bad}} &\equiv \langle \Psi_{SB}^{\text{bad}} | \Psi_{SB}^{\text{bad}} \rangle = e^{-D} \sinh D, \\ P^{\text{good}} &\equiv \langle \Psi_{SB}^{\text{good}} | \Psi_{SB}^{\text{good}} \rangle = e^{-D} \cosh D. \end{aligned} \quad (80)$$

If T is the elapsed time between the preparation and measurement of the qubit, then

$$D = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \tilde{\Delta}(\omega) \frac{4 \sin^2(\omega T/2)}{\omega^2}, \quad (81)$$

where

$$\Delta(t-s) \equiv \langle \phi(t) \phi(s) \rangle = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-i\omega(t-s)} \tilde{\Delta}(\omega) \quad (82)$$

(we have assumed the noise is stationary).

In the case of zero-temperature Ohmic noise, with $\tilde{\Delta}(\omega)$ given by Eq. (17), we find that

$$\begin{aligned} D &= - \int_0^T dt \int_0^T ds \frac{A}{(t-s-i\tau_c)^2} = A \ln \left(\frac{T^2 + \tau_c^2}{\tau_c^2} \right) \\ &\approx 2A \ln(T/\tau_c). \end{aligned} \quad (83)$$

Thus the quantity D (an upper bound on the probability P^{bad} of a measurement error) has only a mild logarithmic sensitivity to the ultraviolet cutoff τ_c^{-1} , in contrast to the power dependence on the cutoff found in Eq. (43). This improvement occurs because D is found by integrating the bath two-point function $\Delta(t)$, rather than its modulus $|\Delta(t)|$, which can be justified because the perturbation H_{SB} commutes with the ideal system Hamiltonian H_S . Even this logarithmic dependence on τ_c may be spurious; it arises because we have assumed that the ideal qubit preparation (at time $t=0$) and qubit measurement (at time $t=T$) are instantaneous. The divergence would be softened further if we used a smoother model of preparation and measurement.

Perhaps the logarithmic dependence of the error probability on the elapsed time T should not be taken too seriously; it applies only if the noise spectrum is Ohmic down to a frequency of order T^{-1} . Let us nevertheless pursue the implications of this behavior. The crux of the scheme formulated in [26] is a teleported logical controlled-NOT (CNOT) gate protected against dephasing by an n -qubit repetition code (where n is odd). This CNOT gadget contains four logical measurements, each of which is decoded by a majority vote. Furthermore, for each qubit, there are at most $3n$ time steps (each of duration t_0) in between the preparation and measurement of the qubit, where a CPHASE gate acts on the qubit in each step. Therefore, the probability ϵ_{CNOT} of an encoded error in this CNOT gadget can be bounded as

$$\epsilon_{\text{CNOT}} \leq 4 \binom{n}{\frac{n+1}{2}} (P^{\text{bad}})^{(n+1)/2}, \quad (84)$$

where

$$P^{\text{bad}} \leq D \leq 2A \ln[(3n+2)t_0/\tau_c] \quad (85)$$

(here we have allowed the noise to act for a time t_0 during each CPHASE gate and also during the initial preparation and final measurement). Hence the logical CNOT gate is well protected if A is small and t_0/τ_c is not too large.

While the underlying noise model is Gaussian dephasing noise, under our assumptions the effective noise model for the CNOT gadgets is independent stochastic noise. Just to be specific, suppose that $A=10^{-3}$ and $t_0/\tau_c=10^3$. Then, for code length $n=9$, Eq. (84) yields $\epsilon_{\text{CNOT}} < 1.85 \times 10^{-6}$. This CNOT error rate is well below the accuracy threshold for the local stochastic noise model, indicating that these logical CNOT gates are adequate for scalable quantum computing.

VIII. CONCLUSIONS

The quantum accuracy threshold theorem indicates that scalable quantum computing is feasible in principle. But will fault-tolerant quantum computation *really* work? One concern is that the noise models assumed by theorists are highly idealized, at best crude approximations to the noise in actual devices. In formulating these models, one desires on the one hand to capture essential features of realistic noise, but on the other hand to allow a succinct and elegant analysis of the computation’s reliability.

Seeking an appropriate balance between these two desiderata, we have proved in this paper a version of the threshold theorem that applies to Gaussian quantum noise, which is physically well motivated and analytically tractable. Our result shows that quantum computing is scalable if the noise power spectrum obeys a certain condition. Compared to previous results regarding the effectiveness of fault-tolerant methods against non-Markovian noise [5–7], our threshold condition has two advantages: it is expressed in terms of experimentally observable features of the noise and it is less sensitive to high-frequency noise.

As mentioned in Sec. VI, it might be useful to extend our results by relaxing the noise model in several ways, for example, by including weak non-Gaussian corrections to the bath fluctuations or by modeling more realistically the dissipative flow of heat from system to bath. It should also be possible to make further improvements in the sensitivity of the threshold condition to high-frequency noise; however, an improved condition would be likely to depend on the details of the frequency spectrum of the ideal system dynamics, and deriving it would require a more complicated analysis.

Experimenters tend to worry less about high-frequency noise than about low-frequency noise, particularly $1/f$ dephasing noise. We anticipate that low-frequency noise in quantum gates can be suppressed substantially through clever design of pulse sequences, leaving weak residual noise to be tamed via the fault-tolerant methods we have studied here. Joining pulse shaping methods with fault-tolerant circuit construction will be a fruitful topic for future research.

ACKNOWLEDGMENTS

We thank Panos Aliferis, Matt Hastings, Alexei Kitaev, Eduardo Novais, and Gil Refael for useful discussions. This research is supported in part by DOE under Grant No. DE-FG03-92-ER40701, NSF under Grant No. PHY-0456720, NSA under ARO Contract No. W911NF-05-1-0294, and by the Gordon and Betty Moore Foundation.

APPENDIX: THRESHOLD THEOREM FOR LOCAL NOISE

Here we will briefly sketch the proof of the quantum accuracy threshold theorem for local noise, following the argument in [6].

We assume that the joint evolution of the quantum computer (the system S) and its environment (the bath B) is

governed by the Hamiltonian $H=H_S+H_B+H_{SB}$, where the perturbation H_{SB} is responsible for the deviation of the system from its ideal evolution. Although this framework can be generalized (as discussed in Sec. VI), we also assume that the system qubits are initialized ideally in the pure state $|\Psi_S^0\rangle$ before the Hamiltonian evolution begins and measured ideally after it ends. Furthermore, the initial state of the bath is the pure state $|\Psi_B^0\rangle$. Just before the ideal measurements are performed on the system qubits, the joint state of the system and bath is $|\Psi_{SB}\rangle=U_{SB}|\Psi_{SB}^0\rangle$, where U_{SB} is the time-evolution operator determined by H , and $|\Psi_{SB}^0\rangle=|\Psi_S^0\rangle\otimes|\Psi_B^0\rangle$. We obtain a fault-path expansion for $|\Psi_{SB}\rangle$ by expanding U_{SB} as a power series in H_{SB} , and for each term in this expansion we declare a level-0 circuit location to be bad if H_{SB} acts nontrivially somewhere within that location.

For any specified set \mathcal{I}_r of r locations in the circuit, we denote by $|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r)\rangle$ the sum of all the terms in the fault-path expansion of $|\Psi_{SB}\rangle$ such that all of these r locations are bad. The noise is local with strength ϵ if

$$\| |\Psi_{SB}^{\text{bad}}(\mathcal{I}_r)\rangle \| \leq \epsilon^r. \tag{A1}$$

Our objective is to show that scalable quantum computing is possible provided that $\epsilon < \epsilon_0$, where ϵ_0 is (a lower bound on) the accuracy threshold.

Suppose that a universal set of fault-tolerant level-1 gadgets can be constructed such that a 1-gadget containing fewer than s faulty level-0 gates simulates the corresponding ideal 0-gate correctly. We can estimate the effective noise strength for a level-1 simulation using the following observation: consider a set \mathcal{I} of level-0 locations in a quantum circuit. Then the sum of all fault paths such that at least s of the locations in the set \mathcal{I} are faulty can be expressed as

$$\begin{aligned} & |\Psi_{SB}(\geq s \text{ faults in } \mathcal{I})\rangle \\ &= \sum_{\ell=s}^{|\mathcal{I}|} (-1)^{\ell-s} \binom{\ell-1}{s-1} \left(\sum_{\mathcal{I}_\ell \subseteq \mathcal{I}} |\Psi_{SB}^{\text{bad}}(\mathcal{I}_\ell)\rangle \right), \end{aligned} \tag{A2}$$

where $\sum_{\mathcal{I}_\ell}$ denotes the sum over all subsets of \mathcal{I} that contain ℓ elements. Equation (A2) follows from the “inclusion-exclusion principle” of combinatorics. For example, in the case $s=1$ it becomes

$$\begin{aligned} & |\Psi_{SB}(\geq 1 \text{ fault in } \mathcal{I})\rangle \\ &= \sum_{\mathcal{I}_1 \subseteq \mathcal{I}} |\Psi_{SB}^{\text{bad}}(\mathcal{I}_1)\rangle - \sum_{\mathcal{I}_2 \subseteq \mathcal{I}} |\Psi_{SB}^{\text{bad}}(\mathcal{I}_2)\rangle + \sum_{\mathcal{I}_3 \subseteq \mathcal{I}} |\Psi_{SB}^{\text{bad}}(\mathcal{I}_3)\rangle \\ &\quad - \sum_{\mathcal{I}_4 \subseteq \mathcal{I}} |\Psi_{SB}^{\text{bad}}(\mathcal{I}_4)\rangle + \dots, \end{aligned} \tag{A3}$$

whose origin is easy to understand. The first term counts correctly each fault path with exactly one fault in \mathcal{I} , but it double counts each fault path with exactly two faults, and this overcounting is corrected by the second term. The first term counts three times each fault path with exactly three faults, and the second term subtracts these fault paths $\binom{3}{2}$ times; this undercounting is corrected by the third term, and so on.

The norm of the left-hand side of Eq. (A2) is bounded above by the sum of the norms of the terms on the right-hand side. Using $\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_\ell)\| \leq \varepsilon^\ell$ and denoting $|\mathcal{I}|=A$ we find

$$\begin{aligned} & \| |\Psi_{SB}(\geq s \text{ faults in } A \text{ locations})\rangle \| \\ & \leq \sum_{\ell=s}^A \binom{\ell-1}{s-1} \binom{A}{\ell} \varepsilon^\ell = \binom{A}{s} \varepsilon^s \sum_{\ell=s}^A \binom{A-\ell}{\ell-s} \varepsilon^{\ell-s} \\ & \leq \binom{A}{s} \varepsilon^s \sum_{t=0}^{\infty} \frac{(A-s)^t \varepsilon^t}{t!} = \binom{A}{s} e^{(A-s)\varepsilon} \varepsilon^s \leq \zeta \binom{A}{s} \varepsilon^s. \end{aligned} \quad (\text{A4})$$

Here ζ is a constant satisfying $\zeta \geq e^{(A-s)\varepsilon}$ for values of ε in some specified range of interest, which typically can be chosen such that ζ is close to 1.

In a level-1 simulation of a quantum circuit, let us say that a level-1 gadget is bad if it contains s or more bad level-0 gates and let $\mathcal{I}_r^{(1)}$ denote a set of r specified level-1 gadgets. We assume, for now, that these level-1 gadgets are nonoverlapping, i.e., that no 0-gate is contained in two different 1-gadgets. We denote by $|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(1)})\rangle$ the sum of all terms in the perturbation expansion of $|\Psi_{SB}\rangle$ such that all of the r 1-gadgets in $\mathcal{I}_r^{(1)}$ are bad. By performing an ‘‘inclusion-exclusion’’ sum independently inside each 1-gadget, we find that

$$|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(1)})\rangle = \sum_{\ell_1=s}^{|\mathcal{I}(1)|} (-1)^{\ell_1-s} \binom{\ell_1-1}{s-1} \cdots \sum_{\ell_r=s}^{|\mathcal{I}(r)|} (-1)^{\ell_r-s} \binom{\ell_r-1}{s-1} \left(\sum_{\mathcal{I}(1)_{\ell_1} \subseteq \mathcal{I}(1)} \cdots \sum_{\mathcal{I}(r)_{\ell_r} \subseteq \mathcal{I}(r)} |\Psi_{SB}^{\text{bad}}(\mathcal{I}(1)_{\ell_1} \cup \cdots \cup \mathcal{I}(r)_{\ell_r})\rangle \right), \quad (\text{A5})$$

where $\mathcal{I}(j)$ denotes the set of 0-gates inside the 1-gadget j , for $j \in \{1, 2, 3, \dots, r\}$, and $\sum_{\mathcal{I}(j)_{\ell_j}}$ denotes the sum over all subsets of $\mathcal{I}(j)$ that contain ℓ_j elements.

The local noise condition $\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_\ell)\| \leq \varepsilon^\ell$ implies that

$$\|\Psi_{SB}^{\text{bad}}(\mathcal{I}(1)_{\ell_1} \cup \cdots \cup \mathcal{I}(r)_{\ell_r})\| \leq \prod_{j=1}^r \varepsilon^{\ell_j}. \quad (\text{A6})$$

As above, we can bound the norm of the left-hand side of Eq. (A5) by the sum of the norms of the terms on the right-hand side. Using Eq. (A6), this upper bound factorizes into a product of r sums, each of which can be bounded as in Eq. (A4). We obtain

$$\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(1)})\| \leq \prod_{j=1}^r \varepsilon_j^{(1)}, \quad (\text{A7})$$

where

$$\varepsilon_j^{(1)} = \zeta_j \binom{A_j}{s} \varepsilon^s, \quad (\text{A8})$$

and hence

$$\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(1)})\| \leq (\varepsilon^{(1)})^r, \quad (\text{A9})$$

where

$$\varepsilon^{(1)} = \max_j (\varepsilon_j^{(1)}). \quad (\text{A10})$$

Here $A_j = |\mathcal{I}(j)|$, $\zeta_j \geq \exp[(A_j - s)\varepsilon]$, and the maximum is over all 1-gadgets in the circuit. We can regard $\varepsilon^{(1)}$ as an effective noise strength for the level-1 circuit, which is conveniently expressed in the form

$$\varepsilon^{(1)} = \varepsilon_0 (\varepsilon / \varepsilon_0)^s, \quad (\text{A11})$$

where

$$\varepsilon_0 = \min \left[\zeta \binom{A}{s} \right]^{-1/(s-1)}, \quad (\text{A12})$$

and the minimum is over all 1-gadgets in our universal set.

Now let us say that a level- k gadget is bad if it contains s or more bad $(k-1)$ -gadgets. The bound Eq. (A8) on the norm of the sum $|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(1)})\rangle$ over all fault paths that are bad at level 1 is of the same form as the bound Eq. (A1) on the norm of the sum over all fault paths that are bad at level 0, but with a ‘‘renormalized’’ value of the effective noise strength. This means that in a recursive simulation, in which k -gadgets are constructed using the same circuits as 1-gadgets, but with each 0-gate in the 1-gadget replaced by a $(k-1)$ -gadget, we can use the same combinatoric argument again to estimate the effective noise strength at level k . That is, suppose that

$$\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(k-1)})\| \leq (\varepsilon^{(k-1)})^r, \quad (\text{A13})$$

where $\mathcal{I}_r^{(k-1)}$ is any specified set of r $(k-1)$ -gadgets in a level- $(k-1)$ simulation, and $|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(k-1)})\rangle$ denotes the sum of all fault paths such that all r of the $(k-1)$ -gadgets in $\mathcal{I}_r^{(k-1)}$ are bad. Then we may infer that

$$\|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(k)})\| \leq (\varepsilon^{(k)})^r, \quad (\text{A14})$$

where $\mathcal{I}_r^{(k)}$ is any specified set of r k -gadgets in a level- k simulation, $|\Psi_{SB}^{\text{bad}}(\mathcal{I}_r^{(k)})\rangle$ denotes the sum of all fault paths such that all r of the k gadgets in $\mathcal{I}_r^{(k)}$ are bad, and

$$\varepsilon^{(k)} / \varepsilon_0 = (\varepsilon^{(k-1)} / \varepsilon_0)^s = (\varepsilon / \varepsilon_0)^{s^k}. \quad (\text{A15})$$

The fault-path expansion of a level- k simulation with all together L k -gadgets can be expressed as

$$|\Psi_{SB}\rangle = |\Psi_{SB}^{\text{good}}\rangle + |\Psi_{SB}^{\text{bad}}\rangle, \quad (\text{A16})$$

where $|\Psi_{SB}^{\text{good}}\rangle$ is the sum of all fault paths such that every k gadget is good and $|\Psi_{SB}^{\text{bad}}\rangle$ is the sum of all fault paths such that at least one k gadget is bad. Using Eq. (A14), we can apply the $s=1$ case of Eq. (A4) to the circuit of L k -gadgets, finding

$$\| |\Psi_{SB}^{\text{bad}}\rangle \| \leq L \exp[(L-1)\varepsilon^{(k)}] \varepsilon^{(k)}, \quad (\text{A17})$$

which is small for $L\varepsilon^{(k)} \ll 1$. Furthermore, the arguments in [6] show that if $|\Psi_{SB}\rangle$ is good then the probability distribution $p^{(\text{actual})} = \{p_a^{(\text{actual})}\}$ governing the measurement outcome for the *logical* system qubits (where p_a is the probability of the measurement outcome labeled by a) matches exactly the probability distribution $p^{(\text{ideal})}$ for the measurement outcomes in the ideal computation. Therefore, the deviation of the $p^{(\text{actual})}$ from $p^{(\text{ideal})}$ in the noisy simulation arises only from the small bad component of $|\Psi_{SB}\rangle$; in fact, in the L^1 norm this deviation can be bounded as

$$\delta = \| p^{(\text{actual})} - p^{(\text{ideal})} \|_1 = \sum_a |p_a^{(\text{actual})} - p_a^{(\text{ideal})}| \leq 2 \| |\Psi_{SB}^{\text{bad}}\rangle \|. \quad (\text{A18})$$

Therefore, for $\varepsilon < \varepsilon_0$, the noisy computation becomes highly reliable as the level k of the simulation increases; thus ε_0 is a

lower bound on the accuracy threshold for quantum computation.

In [6], two valuable extensions of this argument were formulated that are useful for pushing the threshold estimate ε_0 higher. First, the argument can be applied to simulations where successive 1-gadgets *overlap*, i.e., have 0-gates in common. By allowing the gadgets to overlap, we can justify the estimate Eq. (A12) when using properly designed gadgets based on a quantum error-correcting code that can correct $s-1$ errors in a code block. Second, we can refine the definition of badness, so that a 1-gadget with s or more faults is declared bad only if the faults occur at a “malignant” set of locations, i.e., only if the 1-gadget processes encoded information incorrectly because of the faults. For example, for gadgets that can correct one error (the $s=2$ case), our estimate of the level-1 effective noise strength improves to

$$\varepsilon^{(1)} = B\varepsilon^2 + D\varepsilon^3 \leq \varepsilon^2/\varepsilon_0, \quad (\text{A19})$$

where B is the number of malignant *pairs* of fault locations in the 1-gadget (maximized over all 1-gadgets), $D\varepsilon^3$ is a correction arising from summing contributions from fault paths with three or more faults in the 1-gadget, and

$$\varepsilon_0^{-1} = \frac{1}{2}B(1 + \sqrt{1 + 4D/B^2}) \quad (\text{A20})$$

is our improved threshold estimate, found by solving the equation $B\varepsilon_0^2 + D\varepsilon_0^3 = \varepsilon_0$.

-
- [1] P. Shor, in *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, Los Alamitos, CA (IEEE Computer Society Press, Los Alamitos, CA, 1996), p. 56.
- [2] D. Aharonov and M. Ben-Or, *Proceedings of the 29th Annual ACM Symposium on Theory of Computing* (ACM, New York, 1998), p. 176; e-print arXiv:quant-ph/9906129.
- [3] A. Yu. Kitaev, *Russ. Math. Surveys* **52**, 1191 (1997).
- [4] E. Knill, R. Laflamme, and W. H. Zurek, *Proc. R. Soc. London, Ser. A* **454**, 365 (1998).
- [5] B. M. Terhal and G. Burkard, *Phys. Rev. A* **71**, 012336 (2005).
- [6] P. Aliferis, D. Gottesman, and J. Preskill, *Quantum Inf. Comp.* **6**, 97 (2006).
- [7] D. Aharonov, A. Kitaev, and J. Preskill, *Phys. Rev. Lett.* **96**, 050504 (2006).
- [8] P. Aliferis, Ph.D. thesis, California Institute of Technology, Pasadena, CA, 2007.
- [9] P. Aliferis and J. Preskill, *Phys. Rev. A* **79**, 012332 (2009).
- [10] P. Aliferis, D. Gottesman, and J. Preskill, *Quantum Inf. Comp.* **8**, 181 (2008).
- [11] B. W. Reichardt, Ph.D. thesis, University of California–Berkeley, CA, 2006.
- [12] E. Knill, *Nature (London)* **434**, 39 (2005).
- [13] R. Raussendorf, J. Harrington, and K. Goyal, *New J. Phys.* **9**, 199 (2007).
- [14] D. P. DiVincenzo and P. Aliferis, *Phys. Rev. Lett.* **98**, 020501 (2007).
- [15] P. Aliferis and B. Terhal, *Quantum Inf. Comp.* **7**, 139 (2007).
- [16] K. M. Svore, D. P. DiVincenzo, and B. M. Terhal, *Quantum Inf. Comp.* **7**, 297 (2007).
- [17] A. Kay, *Phys. Rev. A* **77**, 052319 (2008).
- [18] M. B. Plenio and S. Virmani, e-print arXiv:0810.4340.
- [19] R. Alicki, e-print arXiv:quant-ph/0702050.
- [20] A. P. Hines and P. C. E. Stamp, *Can. J. Phys.* **86**, 541 (2008).
- [21] R. Alicki, M. Horodecki, P. Horodecki, and R. Horodecki, *Phys. Rev. A* **65**, 062101 (2002).
- [22] A. O. Caldeira and A. J. Leggett, *Ann. Phys.* **149**, 374 (1983).
- [23] E. Novais, E. R. Mucciolo, and H. U. Baranger, *Phys. Rev. Lett.* **98**, 040501 (2007).
- [24] E. Novais, E. R. Mucciolo, and H. U. Baranger, *Phys. Rev. A* **78**, 012314 (2008).
- [25] P. Aliferis and A. W. Cross, *Phys. Rev. Lett.* **98**, 220502 (2007).
- [26] P. Aliferis and J. Preskill, *Phys. Rev. A* **78**, 052331 (2008).
- [27] P. Aliferis, F. Brito, D. P. DiVincenzo, J. Preskill, M. Steffen, and B. M. Terhal, *New J. Phys.* **11**, 013061 (2009).