

Fault-tolerant quantum computation against biased noise

Panos Aliferis¹ and John Preskill²¹IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598, USA²Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA

(Received 29 September 2008; published 19 November 2008)

We formulate a scheme for fault-tolerant quantum computation that works effectively against highly biased noise, where dephasing is far stronger than all other types of noise. In our scheme, the fundamental operations performed by the quantum computer are single-qubit preparations, single-qubit measurements, and conditional-phase (CPHASE) gates, where the noise in the CPHASE gates is biased. We show that the accuracy threshold for quantum computation can be improved by exploiting this noise asymmetry; e.g., if dephasing dominates all other types of noise in the CPHASE gates by four orders of magnitude, we find a rigorous lower bound on the accuracy threshold higher by a factor of 5 than for the case of unbiased noise.

DOI: [10.1103/PhysRevA.78.052331](https://doi.org/10.1103/PhysRevA.78.052331)

PACS number(s): 03.67.Pp

Our confidence that large-scale quantum computers can be realized has been boosted by the theory of fault-tolerant quantum computation [1], which establishes that noisy quantum computers can operate reliably if the noise is not too strong. In a fault-tolerant simulation of a quantum circuit, logical qubits processed by the computer are protected from damage using a quantum code and encoded operations are realized by *gadgets* constructed from the computer's fundamental operations; aside from performing the desired transformation on the encoded quantum information, gadgets also exploit the redundancy of the code to correct errors caused by the noise.

Most work on fault-tolerant quantum computation has focused on the design of gadgets that work effectively for generic noise without any special structure. But actually, in many physical settings the noise is expected to be highly biased. If the computational basis $\{|0\rangle, |1\rangle\}$ coincides with the energy-eigenstate basis for the unperturbed qubit, then typically dephasing (loss of phase coherence in the computational basis, due to entanglement with the environment) is far stronger than relaxation (bit flips in the computational basis, due to energy exchange with the environment). While dephasing arises from low-frequency noise, relaxation is dominated by noise whose frequency is comparable to the energy splitting. Typically, this higher-frequency noise has a different physical origin than the low-frequency noise responsible for dephasing, and it can be orders of magnitude weaker. In this paper, we analyze fault-tolerant gadgets that are designed to exploit this bias.

The fault-tolerant scheme we propose is founded on the assumption that the quantum-computing hardware can execute a conditional-phase (CPHASE) gate with highly biased noise, where CPHASE is the diagonal two-qubit gate with eigenvalues $(1, 1, 1, -1)$ in the computational basis. The complete set of fundamental operations performed by our quantum computer is

$$\mathcal{G}_{\text{fund}} = \{\text{CPHASE}, \mathcal{P}_{|+\rangle}, \mathcal{M}_{\sigma_x}\} \cup \{\mathcal{P}_{|+i\rangle}, \mathcal{P}_{|T\rangle}\}, \quad (1)$$

where \mathcal{M}_{σ_x} denotes the measurement of the Pauli operator σ_x and $\mathcal{P}_{|\psi\rangle}$ denotes the preparation of a single qubit in the state $|\psi\rangle$. To construct fault-tolerant \mathcal{G}_{CSS} operations (see be-

low), we will need to prepare the state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and for fault-tolerant universal quantum computation, we will also need to prepare the states $|+i\rangle = \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$ and $|T\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\pi/4}|1\rangle)$. We have not listed the identify operation, which is implicitly applied whenever a qubit is idle.

Our central assumption, that the noise in CPHASE gates is dominated by dephasing, may apply to some proposed gate implementations using semiconductor spins [2] and superconducting circuits [3]; it may also apply to trapped-ion qubits if the CPHASE gates are driven by microwave fields rather than laser pulses [4]. Furthermore, noise in the preparation of the state $|+\rangle$ is trivially “biased” because $|+\rangle$ is an eigenstate of σ_x and noise in the destructive measurement of σ_x has no specific structure because the measurement has a classical output. We show that, through appropriate gadget design, this noise bias can be exploited to improve the accuracy threshold for quantum computation; e.g., assuming that dephasing dominates all other types of noise by four orders of magnitude, we find that the provable accuracy threshold is higher by more than a factor of 5 than for the case of unbiased noise.

Other authors [5–7] have proposed fault-tolerant gadgets for biased noise, but these previous constructions work only if the noise is dominated by dephasing even for some gates that do not commute with σ_z , such as the controlled-NOT (CNOT) gate. In our view, this assumption is not physically well motivated. A biased noise model should be applicable if, during the execution of a gate, the perturbation responsible for the noise couples predominantly to the σ_z components of the qubits. But during the execution of a gate that does not commute with σ_z , the perturbation may not have this property—e.g., if the gate is a single-qubit rotation about the x axis, then to take into account a possible overrotation or underrotation of the qubit we should include a perturbation proportional to σ_x rather than σ_z . And even if the perturbation is dominated by a σ_z term, an insertion of the perturbation *during* the execution of a rotation about the x axis will propagate to a linear combination of σ_z and σ_y , which cannot be described as dephasing noise alone.

Although the biased noise model has a natural formulation in terms of a Hamiltonian that couples the computer to its environment, we will study a *stochastic* version of the

model. A stochastic noise model assigns a probability to each *fault path*—i.e., to each possible set of faulty fundamental operations in the circuit. We speak of *local stochastic noise* with strength ε if, for any r specified fundamental operations in the circuit, the sum of the probabilities of all fault paths with faults at those r locations is no larger than ε^r [8]. In this model no further restrictions are imposed on the noise and, in particular, the trace-preserving quantum operation applied at the faulty locations of each fault path is arbitrary and can be chosen adversarially. Therefore, although ε quantifies the strength of the noise, the faults can be correlated both temporally and spatially. It was shown recently in [9] that an ideal quantum circuit can be simulated accurately and with reasonable overhead provided that ε is smaller than $\varepsilon_{\text{th}} \geq 0.67 \times 10^{-3}$; this rigorous lower bound on the threshold is the best established so far for this noise model.

The noise model that we will analyze in this paper is a refinement of local stochastic noise that admits two different values of the noise strength: ε , quantifying the rate for faults in preparations and measurements and *dephasing* faults in CPHASE gates, and $\varepsilon' \ll \varepsilon$, quantifying the rate for all other types of faults in CPHASE gates. In this model, a fault path indicates not only which locations are faulty, but also, for each CPHASE gate, whether a dephasing fault or some other type of fault has occurred. We speak of *local stochastic biased noise* if the sum of the probabilities of all fault paths that are faulty at r specified locations, of which s are non-dephasing faults at CPHASE gates, is no larger than $\varepsilon^{r-s}(\varepsilon')^s$. For dephasing faults, all Kraus operators are assumed to be diagonal in the computational basis, and for all other types of faults, the Kraus operators are arbitrary. We refer to the ratio ε/ε' as the noise “bias.”

Our scheme for fault-tolerant quantum computation will be protected by a code $\mathcal{C}_1 \triangleright \mathcal{C}_2$ where \triangleright denotes code concatenation. The inner code \mathcal{C}_1 protects against dephasing, mapping highly biased noise to a more balanced *effective* noise model with reduced noise strength. The code \mathcal{C}_2 maps unstructured noise with strength below ε_{th} to noise with negligible strength. We take \mathcal{C}_1 to be a length- n repetition code in the dual basis, where n is odd; the $n-1$ check operators are $I^{\otimes j} \otimes \sigma_x \otimes \sigma_x \otimes I^{\otimes n-j-2}$ ($j=0, 1, \dots, n-2$), and the logical Pauli operators acting on the one encoded qubit are $\sigma_x^L = \sigma_x \otimes I \otimes I \otimes \dots \otimes I$ and $\sigma_z^L = \sigma_z \otimes \sigma_z \otimes \dots \otimes \sigma_z$. The code \mathcal{C}_1 can correct $(n-1)/2 \sigma_z$ errors but provides no protection against σ_x errors. We take \mathcal{C}_2 to be a concatenated CSS code [8]. For a CSS code, the fault-tolerant encoded versions of operations in the set

$$\mathcal{G}_{\text{CSS}} = \{\text{CNOT}, \mathcal{P}_{|0\rangle}, \mathcal{P}_{|+\rangle}, \mathcal{M}_{\sigma_z}, \mathcal{M}_{\sigma_x}\} \quad (2)$$

can be built out of operations that are contained in this set; furthermore, \mathcal{G}_{CSS} operations suffice for measuring the error syndrome.

We will use the fundamental operations in $\mathcal{G}_{\text{fund}}$ (where the only necessary state preparation is $\mathcal{P}_{|+\rangle}$) to construct \mathcal{G}_{CSS} gadgets protected by \mathcal{C}_1 . Combining with known constructions for CSS codes [10], we obtain \mathcal{G}_{CSS} gadgets protected by $\mathcal{C}_1 \triangleright \mathcal{C}_2$. Finally, CSS operations will be extended to a universal set by appending preparations of the states $|+i\rangle$ and $|T\rangle$; high-fidelity encoded copies of these states can be pre-




FIG. 1. (Color online) Scheme for achieving fault-tolerant universal quantum computation against highly biased noise.

pared by teleporting (“injecting”) into the $\mathcal{C}_1 \triangleright \mathcal{C}_2$ block and then performing state distillation [11]. Our scheme for achieving fault-tolerant universal quantum computation is illustrated in Fig. 1 (we denote an operation \mathcal{O} encoded in \mathcal{C}_1 or $\mathcal{C}_1 \triangleright \mathcal{C}_2$ as \mathcal{O}^L or $\bar{\mathcal{O}}$, respectively).

The crux of our gadget constructions, and the basis for our threshold analysis, is the design of the \mathcal{C}_1 -protected CNOT gadget using the operations in $\mathcal{G}_{\text{fund}}$. The key idea is to use a variant of state teleportation that simultaneously executes the encoded gate and extracts the error syndrome. But first, let us discuss how to construct \mathcal{C}_1 -protected gadgets for the other operations in \mathcal{G}_{CSS} .

The destructive measurement of σ_x^L is performed by measuring σ_x for all qubits in the code block and then computing the majority of the measurement outcomes. This measurement is fault tolerant in the following sense: if m of the qubits in the input code block have errors and s of the single-qubit measurements are faulty, then the outcome of the noisy encoded measurement agrees with the outcome of an ideal encoded measurement provided that $m+s \leq (n-1)/2$. The preparation of $|+\rangle_L$ is executed transversally: $\mathcal{P}_{|+\rangle_L} = (\mathcal{P}_{|+\rangle})^{\otimes n}$. This operation is fault tolerant because at least $(n+1)/2$ of the preparations of $|+\rangle$ must be faulty to cause an encoded error.

A nondestructive measurement of σ_z^L is performed with the circuit depicted in Fig. 2: an ancilla qubit is prepared in the state $|+\rangle$, n consecutive CPHASE gates are executed, and then σ_x is measured on the ancilla qubit. If performed only once, this measurement is not fault tolerant, because a single σ_z error acting on the ancilla can flip the outcome; however, fault tolerance can be ensured by repeating the measurement r times, where r is odd, and computing the majority of the outcomes. Although σ_z errors acting on the data qubits do not affect the measurement outcome, they might contribute to a logical σ_z^L error that could affect subsequent operations. It is therefore noteworthy that if the input block has $m\sigma_z$




FIG. 2. (Color online) Gadget that measures σ_z^L (here, $n=3$). The measurement is repeated r times to ensure fault tolerance (here, $r=3$), with the repetitions staggered as shown so that the data qubits are never idle in between consecutive gates.




FIG. 3. (Color online) On the left, a “one-bit teleportation” circuit. On the right, the error correction gadget.

errors and the measurement gadget has s dephasing faults, then there will be no more than $m+s\sigma_z$ errors in the output block.

The preparation of $|0\rangle_L$ is executed by first preparing $|+\rangle_L$ and then performing the nondestructive measurement of σ_z^L . Again, σ_z errors acting on the data qubits will not disturb the eigenvalue of σ_z^L ; only the σ_z errors acting on the ancilla are problematic. Therefore, the fault tolerance of $M_{\sigma_z^L}$ ensures the fault tolerance of $\mathcal{P}_{|0\rangle_L}$. If the measurement result is $\sigma_z^L = -1$, then the prepared state differs from $|0\rangle_L$ by a known logical σ_x^L error. This error need not be corrected; rather, it is used to update the “Pauli frame” of the encoded block [12].

Error correction can be performed by teleporting an encoded block [12]. Because we only need to correct σ_z errors, the error correction gadget can be simplified to an encoded version of the “one-bit teleportation” circuit [13] depicted in Fig. 3. Ideally, the output encoded qubit has the same state as the input encoded qubit, apart from a possible σ_z^L (if the outcome of $M_{\sigma_z^L}$ is -1) and a possible σ_x^L (if the outcome of $M_{\sigma_z^L \sigma_z^L}$ is -1); thus, the Pauli frame is updated based on the measurement outcomes. The nondestructive measurement of $\sigma_z^L \sigma_z^L$ is performed using one ancilla qubit and $2n$ CPHASE gates, and is repeated r times, much as for the measurement of σ_z^L described above. If the input block has $m\sigma_z$ errors and the error-correction gadget has s dephasing faults, then the outcome of $M_{\sigma_x^L}$ agrees with the ideal case, provided $m+s \leq (n-1)/2$. Furthermore, the outcome of $M_{\sigma_z^L \sigma_z^L}$ agrees with an ideal measurement for $s \leq (r-1)/2$, and the number of σ_z errors in the output block is no more than s .

By combining one-bit teleportation gadgets acting on both output blocks with a logical CNOT gate, we obtain the CNOT gadget depicted in Fig. 4, where the upper block is the control block and the lower block is the target block. The measurements of $\sigma_z^L \sigma_z^L$ and $\sigma_z^L \sigma_z^L \sigma_z^L$ are repeated r times using ancilla qubits and CPHASE gates, as described previously. If the input target block has $m_1\sigma_z$ errors, the input control block has $m_2\sigma_z$ errors, and the CNOT gadget contains s dephasing faults, then both $M_{\sigma_x^L}$'s agree with the ideal case provided that $m_1+s \leq (n-1)/2$ and $m_2+s \leq (n-1)/2$; furthermore, each output block contains no more than $s\sigma_z$ errors, and $M_{\sigma_z^L \sigma_z^L}$ and $M_{\sigma_z^L \sigma_z^L \sigma_z^L}$ agree with the ideal measurements for $s \leq (r-1)/2$. Further properties of the CNOT gadget are discussed in Appendix A.

Now consider a circuit constructed from these \mathcal{C}_1 -protected \mathcal{G}_{CSS} gadgets—how accurately does it simulate an ideal circuit? A gadget operates correctly if all of the encoded measurements it contains agree with the ideal case (the case in which the input blocks have no errors and the

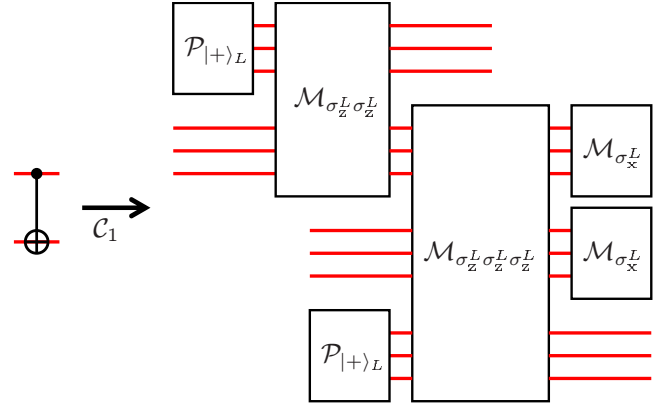


FIG. 4. (Color online) Fault-tolerant CNOT gadget. Pauli operators that update the Pauli frame (not shown) are determined by the measurement outcomes.

gadget contains no faults); otherwise, the gadget fails. For each \mathcal{G}_{CSS} gadget, we have derived an upper bound on its probability of failure in terms of the noise strength ε and the bias factor ε/ε' of the local stochastic biased noise model. See Appendix B for details of this combinatorial analysis.

The largest of these upper bounds (found for the CNOT gadget) is denoted $\varepsilon^{(1)}$; it can be regarded as the effective noise strength of a local stochastic noise model that characterizes the noise in \mathcal{C}_1 -protected \mathcal{G}_{CSS} circuits. In Fig. 5, we have plotted $\varepsilon^{(1)}$ as a function of ε for two different values of the bias. If $\varepsilon^{(1)}$ is below the previously established lower bound on the accuracy threshold for CSS operations ($\varepsilon_{\text{th}}^{\text{CSS}} \geq 0.67 \times 10^{-3}$ [9]), then we can choose the code \mathcal{C}_2 so that the \mathcal{G}_{CSS} gadgets protected by $\mathcal{C}_1 \triangleright \mathcal{C}_2$ are arbitrarily accurate. Thus we set $\varepsilon^{(1)} = 0.67 \times 10^{-3}$ and choose n so that ε is as large as possible. If the bias is 10^4 , then the maximum value is $\varepsilon = 2.50 \times 10^{-3}$, occurring at $n=r=11$.

Furthermore, as we show in Appendix C, the injection and distillation of the $|+i\rangle$ and $|T\rangle$ states can be performed effectively for $\varepsilon \leq 2.50 \times 10^{-3}$. We conclude that 2.50×10^{-3} is a lower bound on the accuracy threshold for universal quantum computation under local stochastic biased noise with

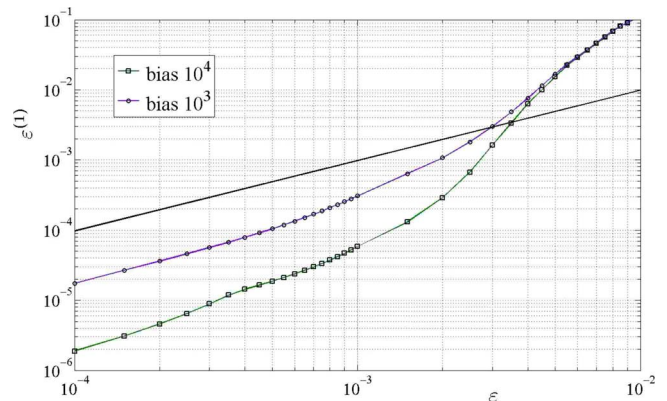


FIG. 5. (Color online) Upper bounds on the effective noise strength $\varepsilon^{(1)}$ for \mathcal{C}_1 -protected \mathcal{G}_{CSS} operations as a function of ε and the bias ε/ε' (where for each value of ε , we optimize over n and r). The straight line with slope unity serves as a guide to the eye.

bias 10^4 . This is an improvement by about a factor of 4 compared to the case of unstructured noise. On the other hand, the improvement is more modest for smaller values of bias; e.g., for a bias of 10^3 , our lower bound becomes 1.54×10^{-3} (where the maximum value $\varepsilon = 1.54 \times 10^{-3}$ occurs at $n=r=7$), an improvement by about a factor of 2.

These results can be further improved by modifying the decoding procedure for $\mathcal{C}_1 \triangleright \mathcal{C}_2$. Recall that the outcome of each \mathcal{C}_1 -protected measurement is determined by a majority vote, and note that this outcome is more likely to be wrong if the vote is “close”—i.e., if the majority has just one more vote than the minority. For example, if $(n-1)/2$ of the σ_x measurements inside $\mathcal{M}_{\sigma_x^L}$ disagree with the choice of the majority, then there might be $(n+1)/2$ errors in the block, resulting in an encoded error. But if only $(n-3)/2$ qubits disagree with the majority, then there must be at least $(n+3)/2$ errors to cause an encoded error. This observation also applies to the majority voting in $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ and $\mathcal{M}_{\sigma_z^L \sigma_z^L \sigma_z^L}$.

The code \mathcal{C}_2 can be decoded more reliably by exploiting information concerning which \mathcal{C}_1 -protected measurements have close votes; see Appendix D for details. Using this more sophisticated decoding method, we find that the accuracy threshold improves to 2.09×10^{-3} for bias 10^3 and improves to 3.51×10^{-3} for bias 10^4 .

Fault-tolerant methods will be essential for achieving large-scale quantum computation. These methods can be more effective when customized for the particular properties of the noise in the computing hardware. In this paper, we have explained how to exploit noise asymmetry in diagonal gates to make fault-tolerant quantum computing work better. We have analyzed the performance of our scheme for local stochastic biased noise; using techniques described in [8,14,15], a more realistic Hamiltonian biased noise model could also be analyzed.

A notable property of our constructions is that the only fundamental operation used by the quantum computer, other than single-qubit preparations and measurements, is the two-qubit CPHASE gate. This feature is attractive, because in some physical settings CPHASE gates are particularly easy to execute with highly biased noise and reasonable fidelity; e.g., our companion paper [16] discusses how the encoding scheme we have formulated here applies to superconducting flux qubits.

We thank David DiVincenzo, Daniel Gottesman, and Gabriel Mendoza for useful discussions. This research is supported in part by DOE under Grant No. DE-FG03-92-ER40701, NSF under Grant No. PHY-0456720, and NSA under ARO Contract No. W911NF-05-1-0294.

APPENDIX A: COMMENTS ON THE CNOT GADGET

In our discussion in the main text, the measurement of $\sigma_z^L \sigma_z^L$ is repeated r times inside $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ and the measurement of $\sigma_z^L \sigma_z^L \sigma_z^L$ is also repeated r times inside $\mathcal{M}_{\sigma_z^L \sigma_z^L \sigma_z^L}$. In general, the number of repetitions could be different in the two cases r_1 and r_2 for the measurements of $\sigma_z^L \sigma_z^L$ and $\sigma_z^L \sigma_z^L \sigma_z^L$, respectively, and for now we will distinguish r_1 and r_2 from n so that the counting we describe below will be more trans-




FIG. 6. (Color online) The full circuit of the CNOT gadget (here, for compactness we have chosen $r_1=r_2=n=2$, but the structure of the circuit is similar for odd values).

parent. In fact, later on we will set $r_1=r_2=n$, which turns out to be optimal or nearly optimal in the cases we have studied. A further advantage of the choice $r_1=r_2=n$ is that we can eliminate storage locations (where qubits are idle) in the CNOT gadget by staggering the measurements as in Fig. 2; for this reason, we will not include any faults at storage locations in our estimate of the failure probability.

A noteworthy property of the CNOT gadget is that, if $r_1=r_2=n$ and if the measurements are staggered as in Fig. 2, then the latest operations on the output blocks act one time step *before* the earliest operations on the input blocks. This property is obscured by the diagrammatic notation in Fig. 4, but it is evident once we consider the full circuit as in Fig. 6. Let us say that a data qubit “interacts” in a time step in which it is coupled to an ancilla qubit by a CPHASE gate. We choose a standard ordering for the n qubits in each block, such that the interactions of qubit j lag $j-1$ time steps behind the interactions of qubit 1. Then, in the CNOT gadget, qubit 1 in the output control block interacts during time steps 1 through n and qubit 1 of the output target block interacts during time steps 1 through n . Meanwhile, qubit 1 in the input control block interacts during time steps $n+1$ through $3n$ and qubit 1 in the input target block interacts during time steps $n+1$ through $2n$. Therefore, in time step $n+1$, as qubit 1 in each input block begins to interact, qubit 1 in each output block is ready for execution of the next gadget. This is a characteristic feature of gate teleportation; it implies that a circuit of \mathcal{G}_{CSS} operations can be simulated in constant depth, independent of the size and depth of the simulated circuit.

APPENDIX B: THE THRESHOLD FOR \mathcal{G}_{CSS} OPERATIONS

Among all \mathcal{G}_{CSS} gadgets, the CNOT gadget contains the largest number of fundamental operations. Therefore, if we derive an upper bound on its probability of failure, this




FIG. 7. (Color online) A CNOT gadget preceded by CNOT gadgets acting on each of its input blocks.

bound will also apply to all the other \mathcal{G}_{CSS} gadgets.

To estimate the probability of failure for a CNOT gadget, we first observe that a σ_x error on a data qubit or σ_z errors acting on many ancilla qubits can cause an incorrect outcome of $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ or $\mathcal{M}_{\sigma_x^L \sigma_x^L}$. The outcome of $\mathcal{M}_{\sigma_x^L}$ may be incorrect due to σ_z errors acting on multiple qubits in a single block or due to a σ_x error acting on an ancilla qubit that propagates repeatedly to generate many σ_z errors in the code block.

In addition, we must take into account errors in preceding gadgets that could propagate into the CNOT gadget we are considering; e.g., Fig. 7 depicts a CNOT gadget preceded by CNOT gadgets acting on each of its input blocks. A σ_z error in one of the preceding gadgets might affect the outcome of an $\mathcal{M}_{\sigma_x^L}$ in the current gadget, and a σ_x error in one of the preceding gadgets can alter the outcome of $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ or $\mathcal{M}_{\sigma_x^L \sigma_x^L}$ in the current gadget.

To understand the effect of σ_x errors in the CNOT gadgets, note that an incorrect outcome of $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ produces a σ_x^L error acting on the CNOT gadget's output control block and an incorrect outcome of $\mathcal{M}_{\sigma_x^L \sigma_x^L}$ produces a σ_z^L error acting on the CNOT gadget's output target block. Consider, e.g., the control block of the final CNOT gadget for the case depicted in Fig. 7, and suppose that a single nondephasing fault in a CPHASE gate contained in the $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ of the immediately preceding gadget alters the outcome of that measurement and also of the $\mathcal{M}_{\sigma_x^L \sigma_x^L}$ in the later gadget. Then this one fault causes logical errors in each of two consecutive gadgets. However, the logical error in the earlier CNOT gadget is a σ_x^L acting on its output target block and has no effect on its output control block. Therefore, we can propagate the logical

error forward from the earlier gadget to the later gadget; i.e., we may just as well say that the earlier gadget is executed properly and the logical error occurs only in the later gadget.

More generally, whenever a single σ_x error causes logical errors in two consecutive gadgets, the logical error in the earlier gadget can be propagated forward into the later gadget in this way. Thus, we may hold the σ_x error responsible for only the failure of the later gadget and we may disregard the damage it inflicts on the earlier gadget. In effect, then, a single nondephasing fault occurring with probability ε' can cause the failure of only one of our \mathcal{C}_1 -protected gadgets.

A measurement of $\sigma_z^L \sigma_z^L$ uses $2n$ CPHASE gates and a measurement of $\sigma_x^L \sigma_x^L$ uses $3n$ CPHASE gates. Therefore, the CNOT gadget contains $(2r_1 + 3r_2)n$ CPHASE gates. We pessimistically assume that any nondephasing fault in a CPHASE gate that is either contained within the CNOT gadget or that propagates into the CNOT gadget causes the gadget to fail. We denote by $\varepsilon_{-d}^{(1)}$ the probability of failure due to a nondephasing fault in a CPHASE gate and conclude that

$$\varepsilon_{-d}^{(1)} \leq (2r_1 + 3r_2 + 2r)n\varepsilon', \quad (\text{B1})$$

where $r \leq \max(r_1, r_2)$. Here, for each input block, rn is an upper bound on the number of CPHASE gates in the preceding gadget where errors can propagate into the current gadget.

Now, suppose that there are no CPHASE gates with nondephasing faults and consider the probability of failure of the CNOT gadget due to dephasing faults at CPHASE gates and due to faults in preparations and measurements. We may assume without loss of generality that the faults at the operations $\mathcal{P}_{|+\rangle}$ and \mathcal{M}_{σ_x} have diagonal Kraus operators, as σ_x errors have no effect at these locations.

A dephasing fault can alter the outcome of a $\sigma_z^L \sigma_z^L$ or $\sigma_x^L \sigma_x^L$ measurement only if it acts on the ancilla qubit used during the measurement. For each of these logical measurements there is one preparation and one measurement operation; furthermore, there are $2n$ CPHASE gates for the $\sigma_z^L \sigma_z^L$ measurement and $3n$ CPHASE gates for the $\sigma_x^L \sigma_x^L$ measurement. We therefore obtain upper bounds on the probability of failure for the majority vote of the repeated measurements:

$$\varepsilon(\mathcal{M}_{\sigma_z^L \sigma_z^L}) \leq \binom{r_1}{\frac{r_1 + 1}{2}} [(2n + 2)\varepsilon]^{(r_1 + 1)/2}, \quad (\text{B2})$$

$$\varepsilon(\mathcal{M}_{\sigma_x^L \sigma_x^L}) \leq \binom{r_2}{\frac{r_2 + 1}{2}} [(3n + 2)\varepsilon]^{(r_2 + 1)/2}. \quad (\text{B3})$$

A measurement of σ_x^L can fail if the majority of the qubits in the block have errors. For each qubit, the error can arise during the preparation of the qubit, the measurement of the qubit, or a CPHASE gate that acts on the qubit. Therefore, an upper bound on the probability of failure is

$$\varepsilon(\mathcal{M}_{\sigma_x^L}^{[1]}) \leq \binom{n}{\frac{n + 1}{2}} [(r + r_1 + r_2 + 2)\varepsilon]^{(n + 1)/2} \quad (\text{B4})$$

for the measurement of the control block and




FIG. 8. (Color online) Teleportation circuits for the Clifford group gates $Q = e^{i(\pi/4)\sigma_x}$ and $S = e^{-i(\pi/4)\sigma_z}$ and for the non-Clifford gate $T = e^{-i(\pi/8)\sigma_z}$. For the Clifford group gates the measurement determines a Pauli operator that updates the Pauli frame, and for the non-Clifford gate the measurement determines a non-Pauli correction that must be applied in the next step.

$$\varepsilon(\mathcal{M}_{\sigma_x}^{[2]}) \leq \left(\frac{n}{n+1} \right) [(r+r_2+2)\varepsilon]^{(n+1)/2} \quad (\text{B5})$$

for the measurement of the target block. Denoting by $\varepsilon_d^{(1)}$ the probability of failure due to faults other than nondephasing faults in CPHASE gates, we obtain

$$\begin{aligned} \varepsilon_d^{(1)} &\leq \varepsilon(\mathcal{M}_{\sigma_z}^{[2]}) + \varepsilon(\mathcal{M}_{\sigma_z}^{[1]}) + \varepsilon(\mathcal{M}_{\sigma_x}^{[1]}) + \varepsilon(\mathcal{M}_{\sigma_x}^{[2]}) \\ &\leq 2 \left(\frac{n}{n+1} \right) [(2n+2)^{(n+1)/2} + (3n+2)^{(n+1)/2}] \varepsilon^{(n+1)/2}, \end{aligned} \quad (\text{B6})$$

where to obtain the second inequality we have substituted $r_1=r_2=r=n$. Our upper bound on the total probability of failure for the CNOT gadget is

$$\varepsilon^{(1)} \leq \varepsilon_{-d}^{(1)} + \varepsilon_d^{(1)} \leq 7n^2 \varepsilon' + \varepsilon_d^{(1)}. \quad (\text{B7})$$

The quantity $\varepsilon^{(1)}$ is the effective noise strength for our \mathcal{C}_1 -protected \mathcal{G}_{CSS} gadgets. In particular, if we set $r=n=11$, $\varepsilon=2.50 \times 10^{-3}$, and $\varepsilon/\varepsilon'=10^4$, we find $\varepsilon^{(1)} < 0.67 \times 10^{-3}$ so that the effective noise strength is below the threshold $\varepsilon_{\text{th}}^{\text{CSS}} \geq 0.67 \times 10^{-3}$ for \mathcal{G}_{CSS} gadgets protected by \mathcal{C}_2 [9]. Thus 2.50×10^{-3} is a lower bound on the accuracy threshold for \mathcal{G}_{CSS} operations assuming a local stochastic biased noise model with bias 10^4 .

APPENDIX C: ACCURACY THRESHOLD FOR UNIVERSAL QUANTUM COMPUTATION

If ε is below the \mathcal{G}_{CSS} threshold, \mathcal{G}_{CSS} gadgets protected by $\mathcal{C}_1 \triangleright \mathcal{C}_2$ are highly reliable. To extend our gadgets to a universal set protected by $\mathcal{C}_1 \triangleright \mathcal{C}_2$, we use gate teleportation as shown in Fig. 8. Provided we can prepare the state $|+i\rangle$,




FIG. 9. (Color online) Injection of the state $|\psi\rangle$ into the $\mathcal{C}_1 \triangleright \mathcal{C}_2$ block by using teleportation. After the encoded Bell state $|\bar{\Phi}_0\rangle$ is prepared, one $\mathcal{C}_1 \triangleright \mathcal{C}_2$ block is decoded to \mathcal{C}_1 . Then, a Bell measurement is performed on the \mathcal{C}_1 block and the input state $|\psi\rangle$.

we can use the \mathcal{G}_{CSS} operations CNOT, $\mathcal{P}_{|+i\rangle}$, \mathcal{M}_{σ_x} , and \mathcal{M}_{σ_z} to teleport $Q = \exp[i(\pi/4)\sigma_x]$ and $S = \exp[-i(\pi/4)\sigma_z]$. Together with the CNOT gate, Q and S suffice to generate the Clifford group. Provided we can prepare the state $|T\rangle$, we can go beyond the Clifford group and achieve universality by using \mathcal{G}_{CSS} operations and S to teleport the gate $T = \exp[-i(\pi/8)\sigma_z]$. Thus, we can do reliable universal quantum computation if we can perform CSS operations reliably and we can also prepare high-fidelity copies of the state $|+i\rangle$ (the eigenstate of σ_y with eigenvalue +1) and the state $|T\rangle$ (the eigenstate of $S\sigma_x$ with eigenvalue +1).

Furthermore, if we are able to prepare noisy copies of $|+i\rangle$ and $|T\rangle$ that are not *too* noisy, then high-fidelity copies can be generated via state distillation protocols [11]. These protocols are based on CSS stabilizer codes, for which \mathcal{G}_{CSS} operations suffice to measure the error syndrome and to decode. The distillation protocol for $|+i\rangle$ uses Steane's [7,1,3] CSS code. In each round of the protocol, the code's check operators are measured for seven noisy copies of the input state; the encoded qubit is accepted and decoded if the error syndrome is trivial. The state $|+i\rangle$ is prepared successfully unless at least three of the input states have errors. Similarly, the distillation protocol for $|T\rangle$ uses a [15,1,3] CSS code. In each round of the protocol, the code's check operators are measured for 15 noisy copies of the input state; the encoded qubit is accepted and decoded if the error syndrome is trivial. Here, too, the state $|T\rangle$ is prepared successfully unless at least three of the input states have errors.

The error threshold for the $|T\rangle$ distillation protocol was estimated in [11], where it was shown that an input error probability as high as 14.1% can be tolerated if each input state is "twirled" by applying $S\sigma_x$ with probability 1/2. The error threshold for $|+i\rangle$ distillation is even higher. Therefore, if \mathcal{G}_{CSS} gadgets protected by $\mathcal{C}_1 \triangleright \mathcal{C}_2$ are reliable and we can also inject input states into the $\mathcal{C}_1 \triangleright \mathcal{C}_2$ block with probability of error below 14.1%; then, we can do reliable universal quantum computation. (By distilling $|+i\rangle$, we can teleport the S gate, enabling us to perform the twirling step in the $|T\rangle$ distillation protocol.)

The state injection is performed by teleportation as in Fig. 9. Let us use $|\bar{\psi}\rangle$ to denote a state encoded in $\mathcal{C}_1 \triangleright \mathcal{C}_2$, to distinguish it from $|\psi\rangle_L$, the state encoded in \mathcal{C}_1 . To inject the single-qubit state $|\psi\rangle$ into the $\mathcal{C}_1 \triangleright \mathcal{C}_2$ block, first the encoded Bell state $|\bar{\Phi}_0\rangle = \frac{1}{\sqrt{2}}(|\bar{0}\rangle|\bar{0}\rangle + |\bar{1}\rangle|\bar{1}\rangle)$ is prepared, and then one of the code blocks is decoded to \mathcal{C}_1 . To complete the telepor-

tation, a joint Bell measurement is performed on the \mathcal{C}_1 block and the unprotected state $|\psi\rangle$. This procedure prepares the state $|\bar{\psi}\rangle$, up to a logical Pauli operator that is known from the outcome of the Bell measurement.

Because $|\bar{\Phi}_0\rangle$ can be prepared by using \mathcal{G}_{CSS} gadgets which are well protected by $\mathcal{C}_1 \triangleright \mathcal{C}_2$ (we start with the encoded state $|\bar{\mathbb{F}}\rangle$ and the encoded state $|\bar{0}\rangle$, and we apply an encoded CNOT gate), we may assume that the preparation of $|\bar{\Phi}_0\rangle$ is flawless. Thus the state injection might fail because of a decoding error, because of an error in the Bell measurement or because of a fault during the single-qubit preparation of $|\psi\rangle$.

Now let us suppose that $\mathcal{C}_2 = \mathcal{C}^{\circ k}$ is obtained by concatenating the CSS code \mathcal{C} all together k times. The decoding of $\mathcal{C}_1 \triangleright \mathcal{C}_2$ is performed recursively: In the first step, $\mathcal{C}_1 \triangleright \mathcal{C}^{\circ k}$ is decoded to $\mathcal{C}_1 \triangleright \mathcal{C}^{\circ(k-1)}$ using gadgets protected by $\mathcal{C}_1 \triangleright \mathcal{C}^{\circ(k-1)}$; then, $\mathcal{C}_1 \triangleright \mathcal{C}^{\circ(k-1)}$ is decoded to $\mathcal{C}_1 \triangleright \mathcal{C}^{\circ(k-2)}$ using gadgets protected by $\mathcal{C}_1 \triangleright \mathcal{C}^{\circ(k-2)}$ and so on. In the last step, $\mathcal{C}_1 \triangleright \mathcal{C}$ is decoded to \mathcal{C}_1 using gadgets protected by \mathcal{C}_1 . Let us denote by $\varepsilon(\mathcal{D})$ the probability that a logical error occurs at any level during this recursive decoding. If decoding is staggered so that no qubits are idle during the Bell measurement, and if the measurement of $\sigma_z^L \sigma_z^L$ is repeated r times inside $\mathcal{M}_{\sigma_z^L \sigma_z^L}$, the probability of a state injection error is

$$\varepsilon(\mathcal{P}_{|\bar{\psi}\rangle}) \leq \varepsilon(\mathcal{D}) + \varepsilon_{\text{BM}} + \varepsilon, \quad (\text{C1})$$

where ε accounts for the probability of a fault in the single-qubit preparation of $|\psi\rangle$ and

$$\begin{aligned} \varepsilon_{\text{BM}} \leq & (2rn + r)\varepsilon' + (1 + r)\varepsilon + \binom{r}{\frac{r+1}{2}} [(n+3)\varepsilon]^{(r+1)/2} \\ & + \binom{n}{\frac{n+1}{2}} [(2r+2)\varepsilon]^{(n+1)/2}. \end{aligned} \quad (\text{C2})$$

Here, $(2rn+r)\varepsilon'$ bounds the probability of error in the Bell measurement due to a nondephasing fault in a CPHASE gate; the fault could occur in one of the $r(n+1)$ gates contained in $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ or in one of the rn gates contained in a measurement in the immediately preceding \mathcal{C}_1 -protected CNOT gadget (which is part of the recursive decoding circuit). Furthermore, $(1+r)\varepsilon$ bounds the probability of error in \mathcal{M}_{σ_x} due to a dephasing fault; the measured qubit participates in r CPHASE gates contained in $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ and also in the measurement itself [the probability of a fault in the preparation of this qubit in the state $|\psi\rangle$ has already been included in Eq. (C1)]. The next to last term bounds the probability of error in $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ due to a dephasing fault; the ancilla qubit in each of the r measurements inside $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ participates in one $\mathcal{P}_{|+\rangle}$, $n+1$ CPHASE gates, and one \mathcal{M}_{σ_x} . Finally, the last term bounds the probability of an error in \mathcal{M}_{σ_x} ; each qubit in the measured block participates in one $\mathcal{P}_{|+\rangle}$, r CPHASE gates contained in the preceding \mathcal{C}_1 -protected CNOT gadget, r CPHASE gates contained in $\mathcal{M}_{\sigma_z^L \sigma_z^L}$, and one \mathcal{M}_{σ_x} .




FIG. 10. (Color online) A Bell state $|\Phi_0\rangle_L$ is prepared by starting with two blocks in the state $|+\rangle_L$ and then measuring $\sigma_z^L \sigma_z^L$.

If we set $r=n=11$, $\varepsilon \leq 2.50 \times 10^{-3}$, and $\varepsilon' \leq 10^{-4}\varepsilon$, we find $\varepsilon_{\text{BM}} \leq 3.01\%$. Since $\varepsilon(\mathcal{D}) \leq 8.24\%$ for $\varepsilon^{(1)} \leq 0.67 \times 10^{-3}$ [9], we conclude that $\varepsilon(\mathcal{P}_{|\bar{\psi}\rangle}) \leq 11.5\%$ which is below the 14.1% distillation threshold. Thus 2.50×10^{-3} is a lower bound on the accuracy threshold for universal quantum computation under local stochastic biased noise model with bias 10^4 .

APPENDIX D: IMPROVED THRESHOLD VIA FLAGGING AND MESSAGE PASSING

We can improve our lower bound on the accuracy threshold by using a more sophisticated decoding procedure for $\mathcal{C}_1 \triangleright \mathcal{C}_2$. We note that the syndrome information for \mathcal{C}_1 is helpful for optimizing the decoding of \mathcal{C}_2 in the concatenated block; yet the decoding procedure that we have described so far makes no use of this information—after \mathcal{C}_1 is decoded, the \mathcal{C}_1 syndrome is discarded. Now we consider a new decoding procedure where some information about the \mathcal{C}_1 syndrome is retained and used in the decoding of \mathcal{C}_2 .

For the sake of clarity, we continue to make a distinction between r_1 , r_2 , and n , even though we will set them equal later on. We say that a vote is “close” if the winners have one more vote than the losers. Thus $\mathcal{M}_{\sigma_x^L}$ has a close vote if the error syndrome indicates $(n-1)/2$ errors in the block, $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ has a close vote if $(r_1-1)/2$ of the measurements disagree with the majority, and $\mathcal{M}_{\sigma_z^L \sigma_z^L \sigma_z^L}$ has a close vote if $(r_2-1)/2$ measurements disagree with the majority. If a gadget contains no close votes, then we decode \mathcal{C}_1 as usual. But if the gadget contains a close vote, then a *flag* is raised after decoding. The flag signifies that the gadget has a higher than usual probability of failure, information that will be exploited during decoding at the next level up in the concatenated block, using a scheme described in [12].

For simplicity, we consider a version of the scheme in [12] where \mathcal{C}_2 is the concatenated 4-qubit code (\mathcal{C}_4) with check operators $\sigma_x^{\otimes 4}$ and $\sigma_z^{\otimes 4}$; this is the case analyzed in [9]. The basic building blocks for the construction of \mathcal{C}_2 -protected \mathcal{G}_{CSS} operations in this scheme are Bell states $|\Phi_0\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle)$ as on the left of Fig. 10 and CNOT gates followed by single-qubit measurements as on the top of Fig. 11. We can then construct encoded versions of these two basic building blocks by using \mathcal{C}_1 -protected gadgets as on the right of Fig. 10 and the bottom of Fig. 11.

Consider how to refine our estimate of the failure probability for $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ in Fig. 10 taking flagging into account. If there is a flag, then $(r_1+1)/2$ faulty measurements might cause $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ to fail, but without a flag at least $(r_1+3)/2$




FIG. 11. (Color online) On the top, the CNOT gates and measurements used to implement the teleportations of \mathcal{C}_2 blocks and subblocks in Fig. 3 in [9]. On the bottom, a \mathcal{C}_1 -protected implementation of the operations on the top. The outcome of a measurement of $\sigma_z^L \sigma_z^L$ or $\sigma_z^L \sigma_z^L \sigma_z^L$ on the bottom corresponds to the outcome of the measurement of σ_z on the top with the same superscript, while the outcomes of the measurements of σ_x^L on the bottom are combined as shown to give the outcomes of the measurements of σ_x on the top.

faulty measurements are required. Therefore, instead of Eq. (B2) we have

$$\varepsilon_f(\mathcal{M}_{\sigma_z^L \sigma_z^L}) \leq \left(\frac{r_1}{r_1 + 1} \right) [(2n + 2)\varepsilon]^{(r_1+1)/2} \quad (D1)$$

for the probability of failure with a flag and

$$\varepsilon_{-f}(\mathcal{M}_{\sigma_z^L \sigma_z^L}) \leq \left(\frac{r_1}{r_1 + 3} \right) [(2n + 2)\varepsilon]^{(r_1+3)/2} \quad (D2)$$

for the probability of failure without a flag. In fact, because the circuit in Fig. 10 prepares an ancilla state, we find it advantageous to repeat this measurement of $\sigma_z^L \sigma_z^L$ a smaller number of times $t < n$ and to post-select on the cases without a flag. In this case, the conditional probability that $\mathcal{M}_{\sigma_z^L \sigma_z^L}$ fails when it is accepted is

$$\varepsilon(\mathcal{M}_{\sigma_z^L \sigma_z^L} | \text{acc}) \leq \frac{\left(\frac{t}{t+3} \right) [(2n+2)\varepsilon]^{(t+3)/2} + 2nt\varepsilon'}{1 - \left(\frac{t}{t-1} \right) [(2n+2)\varepsilon]^{(t-1)/2} - 2nt\varepsilon'} \quad (D3)$$

where $\left(\frac{t}{t-1} \right) [(2n+2)\varepsilon]^{(t-1)/2}$ bounds the probability that a flag is raised and $2nt\varepsilon'$ is the probability of a nondephasing error in one of the CPHASE gates.

We can perform a similar analysis to bound the probability of failure, with and without a flag, for the measurements $\mathcal{M}_{\sigma_z^L \sigma_z^L}$, $\mathcal{M}_{\sigma_z^L \sigma_z^L \sigma_z^L}$, and $\mathcal{M}_{\sigma_x^L}$ in Fig. 11. These bounds can now be plugged into the analysis of the decoding of \mathcal{C}_2 in [9]. For bias $\varepsilon/\varepsilon' = 10^4$, we find that reliable ($\mathcal{C}_1 \triangleright \mathcal{C}_2$)-protected \mathcal{G}_{CSS} operations can be implemented if ε is below 3.51×10^{-3} (where this optimal value is achieved by choosing $r_1 = r_2 = n = 7$ and $t = 5$). In addition, by an analysis similar to the discussion in Appendix C, we find that for $\varepsilon = 3.51 \times 10^{-3}$ the probability of an error in state injection is $\varepsilon(\mathcal{P}_{|\bar{y}\rangle}) \leq 10.4\%$, which is below the 14.1% distillation threshold. Thus 3.51×10^{-3} is our improved lower bound on the accuracy threshold for universal quantum computation under local stochastic biased noise with bias 10^4 .

[1] P. Shor, in *Proceedings of the 37th Annual Symposium on Foundations of Computer Science, Los Alamitos, CA, 1996* (IEEE Computer Society Press, New York, 1996), p. 56.
 [2] J. M. Taylor *et al.*, *Nat. Phys.* **1**, 177 (2005).
 [3] F. Brito *et al.*, *New J. Phys.* **10**, 033027 (2008).
 [4] C. Ospelkaus *et al.*, *Phys. Rev. Lett.* **101**, 090502 (2008),
 [5] I. Gourlay and J. F. Snowdon, *Phys. Rev. A* **62**, 022308

(2000).
 [6] A. M. Stephens *et al.*, *Phys. Rev. A* **77**, 062335 (2008).
 [7] Z. W. E. Evans *et al.*, e-print arXiv:0709.3875.
 [8] P. Aliferis, D. Gottesman, and J. Preskill, *Quantum Inf. Comput.* **6**, 97 (2006).
 [9] P. Aliferis and J. Preskill, e-print arXiv:0809.5063.
 [10] P. Aliferis, D. Gottesman, and J. Preskill, *Quantum Inf. Comput.*

- put. **8**, 181 (2008).
- [11] S. Bravyi and A. Kitaev, Phys. Rev. A **71**, 022316 (2005).
- [12] E. Knill, Nature (London) **434**, 39 (2005).
- [13] X. Zhou, D. W. Leung, and I. L. Chuang, Phys. Rev. A **62**, 052316 (2000).
- [14] B. M. Terhal and G. Burkard, Phys. Rev. A **71**, 012336 (2005).
- [15] D. Aharonov, A. Kitaev, and J. Preskill, Phys. Rev. Lett. **96**, 050504 (2006).
- [16] P. Aliferis *et al.*, New J. Phys. e-print arXiv:0806.0383.