Error Suppression for Arbitrary-Size Black Box Quantum Operations

Gideon Lee,[1] Connor T. Hann,[2,3,4,5,6] Shruti Puri,[2] S. M. Girvin,[3,4] and Liang Jiang[1,5]
[1]Pritzker School of Molecular Engineering, The University of Chicago, Chicago, Illinois 60637, USA
[2]Department of Applied Physics, Yale University, New Haven, Connecticut 06520, USA
[3]Department of Physics, Yale University, New Haven, Connecticut 06511, USA
[4]Yale Quantum Institute, New Haven, Connecticut 06520, USA
[5]AWS Center for Quantum Computing, Pasadena, California 91125, USA
[6]IQIM, California Institute of Technology, Pasadena, California 91125, USA

(Received 13 November 2022; accepted 23 October 2023; published 9 November 2023)

Efficient suppression of errors without full error correction is crucial for applications with noisy intermediate-scale quantum devices. Error mitigation allows us to suppress errors in extracting expectation values without the need for any error correction code, but its applications are limited to estimating expectation values, and cannot provide us with high-fidelity quantum operations acting on arbitrary quantum states. To address this challenge, we propose to use error filtration (EF) for gate-based quantum computation, as a practical error suppression scheme without resorting to full quantum error correction. The result is a general-purpose error suppression protocol where the resources required to suppress errors scale independently of the size of the quantum operation, and does not require any logical encoding of the operation. The protocol provides error suppression whenever an error hierarchy is respected—that is, when the ancillary controlled-SWAP operations are less noisy than the operation to be corrected. We further analyze the application of EF to quantum random access memory, where EF offers hardware-efficient error suppression.

DOI: 10.1103/PhysRevLett.131.190601

Introduction.—One major obstacle to performing meaningful computation on quantum devices is the presence of noise. Canonically, we expect that the theory of fault-tolerant quantum error correction (FTQEC) codes will enable us to scale quantum computers once we have enough qubits and physical error rates fall below a particular threshold [1–3]. However, being in the noisy intermediate-scale quantum (NISQ) era or early fault-tolerance era means that we are limited in the number and base quality of qubits available [4], which prevents us from performing full fault-tolerant quantum computing. Recent work related to suppressing errors on NISQ devices has focused on error mitigation [5]—for instance, zero noise extrapolation [6,7], quasiprobability decomposition, and probabilistic error cancellation [7–11], learning-based methods such as Clifford data regression [12,13] deep learning noise prediction [14], and virtual distillation [15–17]. Such methods allow the user to suppress errors in extracting expectation values with minimal hardware overhead. The success of such methods in the near term motivates the desire to suppress errors in quantum gates beyond expectation values.

One approach to achieve more robust quantum gates is to use error detection techniques. In the near term, one may not want to use the full formalism of quantum error correction (QEC). One promising alternative approach to detect errors without full QEC is error filtration (EF), which was first introduced as a means to stabilize quantum communication [18]. EF does not seek to mitigate errors in expectation values but rather to protect quantum information during noisy communication. In essence, EF multiplexes a single message, and then attempts to detect and discard the parts of this message in which errors have occurred. Up to postselection, one is able to communicate a message over multiple similarly noisy channels with lower error rates than a single noisy channel. Given a single-channel error that goes as ε, EF is able to suppress errors in the fidelity of the communicated message to ε/T, where T is the number of channels in the multiplexing which corresponds to the effective dimension of the ancilla Hilbert space. Because of its ease of implementation, a successful proof-of-principle experiment was quickly carried out [19]. Recent interest in EF has seen a revival in the context of a more general class of schemes communicating over a quantum superposition of trajectories. Such schemes boast a range of exotic and remarkable results, such as perfect quantum communication over zero-capacity channels [20–22]. Separately, [23] also formalized aspects of EF and derived explicit EF fidelities for loss and dephasing channels. However, until now, EF has mostly been studied in the context of suppressing errors in communication, which is restricted to identity operations [24]. Also relevant shortly is the development of biased noise qubits and bias-preserving gates which suffer from an exponentially

smaller likelihood of bit flips than phase flips [25–27]. In QEC, the use of biased-noise architectures gives rise to much higher fault-tolerant thresholds for QEC codes [28,29]. In the context of this work, they boost the performance of a protocol that might otherwise be infeasible in the near future.

In this Letter, we extend EF to the context of gate-based quantum computation (gate-based EF), and show that this provides a low-overhead means to suppress errors for a large class of quantum operations. The result is a general-purpose error detection protocol where the resources required to suppress errors scale independently of the size of the quantum operation. This has the appealing consequence of allowing us to leverage any small number of additional qubits available in a noisy device to suppress errors in large, complex quantum operations. To establish these results, we provide a general quantum circuit design (Fig. 2) that can filter out errors by employing the noisy operations as black boxes. We stress that EF deals with quantum operations, and provides a means to move beyond suppressing expectation values without the full formalism of QEC.

Gate-based EF with one control qubit.—Suppose we are given a black box that imperfectly carries out some ideal unitary U. We model this as a completely positive and trace-preserving map $\mathcal{U}$ comprising Kraus operators K_i, $i = 0, 1, \ldots, R$ [30], where we define K_0 to be the Kraus operator whose normalized action most resembles U. We assume that queries to this black box are already fairly close to the ideal unitary U. We can formalize this notion by assuming that $\|K_0 - U\| \leq \varepsilon \ll 1$. Then we can write $K_0 = U - \varepsilon\xi$ for some suitably normalized operator ξ with $\|\xi\| = 1$. The infidelity goes as $(1-F)_0 \equiv \langle U\psi|\mathcal{U}(\psi)|U\psi\rangle \sim O(\varepsilon)$ for nonunitary errors, while $(1-F)_0 \sim O(\varepsilon^2)$ for unitary errors [31]. Without any further information, or prying the box apart and subjecting every qubit to QEC, how can we suppress the $O(\varepsilon)$ nonunitary error? We argue that one effective way to do so, leveraging a small number of high quality biased-noise qubits, is gate-based EF.

We begin by introducing gate-based EF with only one control qubit. This provides a gentle introduction to gate-based EF in its simplest possible incarnation and demonstrates that with minimal overhead, gate-based EF is able to help us achieve nontrivial error suppression in the near term. One can think about multiplexing in the original EF protocol as quantum communication over a superposition of trajectories [20]. Inspired by this structure, we want to create a superposition of T queries to $\mathcal{U}$. Figure 1 depicts the minimal implementation of gate-based EF with $T = 2$. To create a superposition of calls to the black box, we need three ingredients. The first is entanglement with a single qubit control register that maintains the superposition between calls to $\mathcal{U}$, the second is a memory register, initialized with the desired input state $|\psi\rangle$, to store the

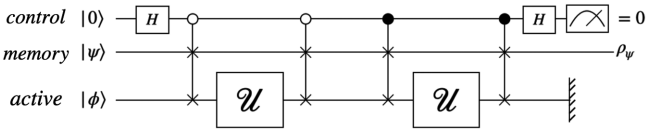

FIG. 1. Gate-based EF with a single control qubit. This circuit makes two calls to an apparatus that implements the noisy process $\mathcal{U}$ as follows: (1) prepare the control register in the equal superposition state $|+\rangle = H|0\rangle$, the intended input state $|\psi\rangle$ in the memory register, and any easy-to-prepare state $|\phi\rangle$ in the active register. (2) Conditioned on the state of the control register being in $|0\rangle$ ($|1\rangle$), the circuit first applies $\mathcal{U}$ to $|\psi\rangle$ ($|\phi\rangle$), then to $|\phi\rangle$ ($|\psi\rangle$). (3) After a final Hadamard transform, the circuit postselects the result conditioned on obtaining the outcome $|0\rangle$. The active register is discarded. This results in the suppression of the infidelity of the apparatus by half [see Eq. (6)].

results of these calls, and the third is an active register, initialized in some (for now) arbitrary state $|\phi\rangle$, for null calls to $\mathcal{U}$, which is discarded at the end of the protocol. We prepare the control register in the equal superposition state $|+\rangle$. Conditioned first on the control register being $|0\rangle$ and then $|1\rangle$, we query $\mathcal{U}$ twice with the input state $|\psi\rangle$. Subsequently, we take a measurement on the control register and postselect on obtaining $|+\rangle$. Finally, we trace out the active register, and are left with the final state ρ_1.

To evaluate the scheme, we use the infidelity

$$(1-F)_1 \equiv 1 - \frac{\langle U\psi|\rho_1|U\psi\rangle}{\mathrm{Tr}\rho_1}, \tag{1}$$

where $|U\psi\rangle = U|\psi\rangle$ is the ideal state we are trying to achieve. While the above infidelity appears state dependent, $(1-F)_1$ as a function of $(1-F)_0$ turns out to be independent of $|\psi\rangle$, $|\phi\rangle$.

First, suppose all the errors in the circuit come from the black box implementation of $\mathcal{U}$. Because of postselection, ρ_1 is not normalized, with $P_1^{(S)} \equiv \mathrm{Tr}\rho_1$ giving the success probability. A straightforward calculation gives the unnormalized state

$$\rho_1 = \frac{1}{2}\mathcal{U}(|\psi\rangle\langle\psi|) + \frac{1}{2}\sum_{i=0,j=0}^{R} K_i|\psi\rangle\langle\psi|K_j^\dagger \mathrm{Tr}\left(\rho_\phi K_i^\dagger K_j\right), \tag{2}$$

and success probability,

$$P_s^{(1)} = \frac{1}{2} + \frac{1}{2}\sum_{i=0,j=0}^{R} \langle\psi|K_j^\dagger K_i|\psi\rangle \mathrm{Tr}\left(\rho_\phi K_i^\dagger K_j\right). \tag{3}$$

Recall our assumption that $K_0 = U - \varepsilon\xi$, which implies that the terms in the sum where $i, j \neq 0$ are $O(\varepsilon^2)$. As such, when computing the fidelity we only have to keep terms where either $i, j = 0$. For clarity of presentation, we now assume the simplest possible model fulfilling our

assumptions, with $K_0 = \sqrt{1-\varepsilon}U, K_1 = \sqrt{\varepsilon}V$, and $V \neq U$ some arbitrary erroneous unitary. We deal with the general case in [31], for which the main result below [Eq. (6)] has exactly the same form. Specializing to this error model, the explicit probability of success is close to 1:

$$1 - P_1^{(S)} \simeq \varepsilon\big(1 - \mathrm{Re}\big\{\langle\psi|U^\dagger V|\psi\rangle \mathrm{Tr}(\rho_\phi V^\dagger U)\big\}\big) \leq 2\varepsilon, \quad (4)$$

where we have discarded $O(\varepsilon^2)$ terms, and the final inequality comes from unitarity of U, V. The inner product with the ideal state is

$$\langle\psi|U^\dagger\rho_1 U|\psi\rangle \simeq \frac{1}{2}F_0 + \frac{1}{2} - \varepsilon\big(1 - \mathrm{Re}\big\{\langle\psi|U^\dagger V|\psi\rangle \mathrm{Tr}(\rho_\phi V^\dagger U)\big\}\big). \quad (5)$$

Inserting Eqs. (4) and (5) into Eq. (1), the infidelity is readily obtained. Up to $O(\varepsilon)$, the V, ϕ dependent terms cancel out, we have

$$(1-F)_1 = \frac{1}{2}(1-F)_0 + O(\varepsilon^2), \quad (6)$$

i.e., the infidelity is halved with a single ancilla qubit. Equation (6) is our first key result, and demonstrates the ability of gate-based EF to suppress errors with minimal overhead comprising a single control qubit and an additional memory register. We highlight that the scaling of $(1-F)_1$ with $(1-F)_0$ is independent of the erroneous unitary V or the active register state $|\phi\rangle$. This latter point suggests that we may initialize the active register in any state, e.g., a thermal state, that is easiest to prepare in the lab. Finally, note that as long as $F_0 > 1/2$, gate-based EF can still provide error suppression even if $\varepsilon \sim O(1)$ [31].

The effects of ancilla noise.—In the near term, we will not have perfectly noiseless ancillae. What happens to the results, Eqs. (4) and (6)? We expect the primary mechanism of introducing additional errors to be the controlled-SWAP (cSWAP) gate. One mitigating factor is the recent discovery of biased-noise cat codes allow us to construct a bias-preserving Toffoli gate [25–27], from which we can construct cSWAP operations in a bias-preserving manner. Recent experimental work [39,40] has demonstrated biases of up to 10^9 can be achieved with cat qubits, and theoretical work [41] demonstrates that such extreme biases can be achieved with modest excitation numbers and low dominant error rates. We thus focus on ancillary phase flip errors, the effects of which are twofold. First, a phase flip error on the control register commutes with the cSWAP operation. Hence, it can be propagated to the end of the circuit, where it is detected by the measurement—i.e., phase flip errors decrease the success probability of the scheme without affecting the fidelity. Assuming this event occurs with some probability p_z over the course of the circuit, we can modify Eq. (4) by substituting $P_1^{(S)} \to P_1^{(S)} - p_z$.

On the other hand, errors on the memory register are more problematic, since they interact with the input state $|\psi\rangle$ in an uncontrolled way. Let p_m be the effective probability of a memory register error during the protocol. Since most near term devices are dominated by gate errors [42], we can assume that these errors come from the cSWAP, so that $p_m \sim p_z$. The worst case assumption that this error completely ruins the query and cannot be detected by post-selection modifies Eq. (6) by $(1-F)_1 \to (1-F)_1 + p_m$. This reveals the ideal operation of gate-based EF to require an error hierarchy. Equation (6) holds when $p_m \ll \varepsilon \ll 1$. However, as long as $p_m < (1-F)_0 \sim \varepsilon$, gate-based EF will still suppress errors if the cSWAP operations are less noisy than $\mathcal{U}$ [31].

Gate-based EF with T control qubits.—Having introduced the base case of gate-based EF applicable to near-term devices, we now generalize the notion of gate-based EF to a situation with $\log T$ control qubits [43]. While the $\log T = 1$ case examined earlier is immediately applicable, we envision the general case to be useful when better ancilla qubits become available. To motivate this, one can imagine having a small number of high quality or error-corrected qubits, which do not suffice for full QEC of the apparatus $\mathcal{U}$. However, applying these qubits as ancillae in gate-based EF, one can still achieve considerable error suppression for $\mathcal{U}$.

The generalization is depicted in Fig. 2. The two modifications to the base case circuit are appending of additional qubits to the control register, and the usage of many cSWAPs. In this case, we have T applications of $\mathcal{U}$ to $|\psi\rangle$ conditioned on the control qubits being in $|00\ldots00\rangle, |00\ldots01\rangle, \ldots, |11\ldots11\rangle$.

When the ancillae are much less noisy than the apparatus, Eq. (6) generalizes to

$$(1-F)_{\log T} \simeq \frac{1}{T}(1-F)_0 + O(\varepsilon^2), \quad (7)$$

independent of $|\psi\rangle$, $|\phi\rangle$. This is a key result of our work, and the full derivation is contained in [31].

To understand the scaling of the success probability with T, we can make the worst-case assumption that the occurrence of an error $K_{i>0}$ on any step causes us to reject the output. This yields the lower bound

$$P_{\log T}^{(S)} \geq 1 - T\varepsilon + O(\varepsilon^2). \quad (8)$$

When $T \ll 1/\varepsilon$, the scheme will still work with high probability, and Eq. (8) allows one to trade off between error suppression and success probability.

Surprisingly, we can often do better than Eq. (8). Under certain favorable conditions, the success probability can be lower bounded by a constant,

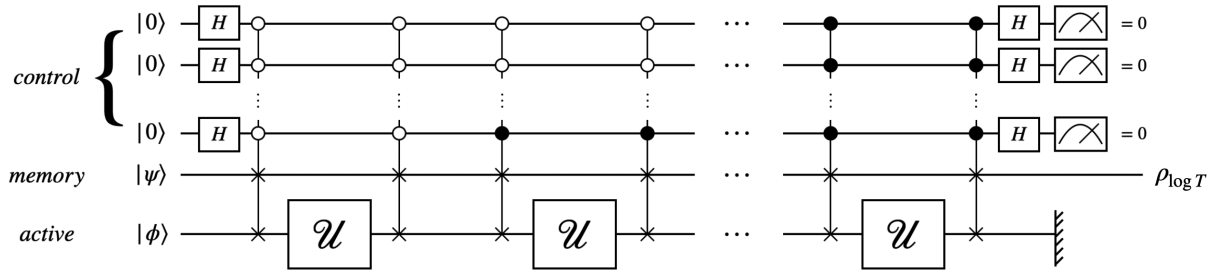

FIG. 2. General circuit for gate-based EF that makes some number T calls to an apparatus that implements the noisy process $\mathcal{U}$. This generalizes the circuit in Fig. 1 by allowing for $\log T$ control qubits, thus making T calls to $\mathcal{U}$ conditioned on each branch $|i\rangle$ of the control register in the computational basis (i.e., $|0\ldots00\rangle, |0\ldots01\rangle, \ldots$).

$$P^{(S)}_{\log T} \geq 1 - 4\varepsilon + \frac{\varepsilon}{T}, \tag{9}$$

which approaches a constant $P^{(S)}_{\log T} \to 1 - 4\varepsilon$ as T increases. These conditions are detailed in [31], but we note a particularly relevant case: if the apparatus $\mathcal{U}$ has biased (Z) noise, one can achieve this bounded success probability by initializing the active register in the $Z = +1$ eigenstate $|00\ldots0\rangle$.

Finally, as long as ancilla errors remain smaller than the apparatus errors, one can still find an optimal point for T where the error suppression of gate-based EF is maximum [31].

Application to QRAM.—To illustrate the practical utility of our scheme, we consider its application to quantum random access memory (QRAM) [44,45]. Error-correcting QRAM has an extremely large hardware overhead. The base hardware overhead of QRAM is already tremendous—in order to prepare a state on $\log N$ address bits, QRAM requires $O(N)$ physical qubits. In data processing applications, relevant values of N could easily reach $N \sim 10^6$–10^9 individually encoded qubits. This problem of hardware overhead is compounded by QRAM being a non-Clifford operation, requiring special techniques [46,47] to implement fault tolerantly [30]. The detailed analysis of [48] corroborates this intuition by demonstrating that a fault-tolerant surface code implementation of QRAM for a memory of size $N \sim 10^6$–10^9 would require some 10^{10}–10^{13} physical qubits.

In contrast to QEC, the resource overhead of gate-based EF scales independently of the size of the desired quantum operation. Additionally, QRAM satisfies the two conditions for the optimal application of gate-based EF. First, the error hierarchy is enforced since N is generally very large, ensuring many more errors occur in the apparatus than the ancillae. Second, one can implement QRAM with biased noise, satisfying the conditions for an upper-bounded failure probability [31]. This ensures that it remains feasible to embed a QRAM with gate-based EF into a quantum algorithm as an oracle. Thus, gate-based EF can suppress QRAM errors in a hardware-efficient manner. For a more complete review of QRAM and the numerical techniques used to simulate it, see [32].

To showcase our scheme's hardware efficiency, we numerically simulate its application to QRAM circuits comprising up to $2^n = 8$ qubits with up to $T = 4$ ancilla qubits in Fig. 3 without ancilla errors. Absent ancilla errors, we see an excellent agreement with the $1/T$ scaling. With ancilla errors [31], we find that there is some T for which gate-based EF is optimal. We simulate the physical qubits of the QRAM with a 0.01 depolarizing error rate per time step, which gives a base infidelity that goes as $O[0.01(\log N)^2]$ [32]. In this regime, we find that our scheme can reduce the query infidelity by over an order of magnitude using only $4 + \log N$ additional qubits. For QRAM, we have shown that one can begin to suppress errors, albeit to a smaller extent than QEC, with only

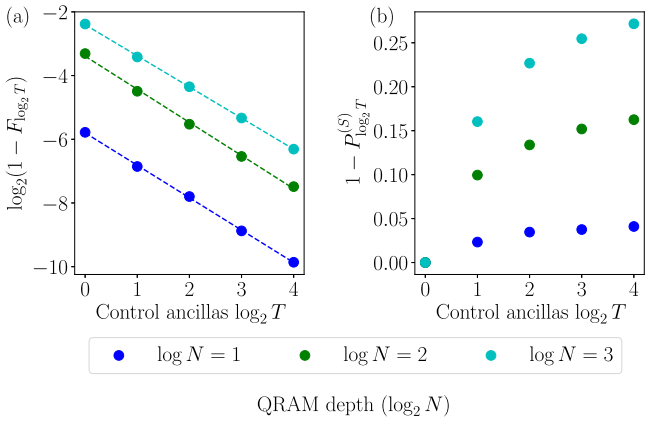

FIG. 3. Gate-based EF applied to QRAM subject to depolarizing errors. (a) Plot of $\log(1 - F_{\log T})$ as a function of $\log T$, where $F_{\log T}$ denotes QRAM query fidelity obtained after gate-based EF with $\log T$ control qubits. The dashed lines indicate linear fits. This demonstrates good agreement with the expected $1/T$ suppression for low apparatus error, with the deviation of the estimated slope from -1 explained by the $O(\varepsilon^2)$ terms for higher apparatus error rates. (b) Plot of failure probabilities $1 - P^{(S)}_{\log T}$ as a function of $\log T$. The simulated failure probability for depth 1,2 QRAMs quickly plateau. All failure probability plots show sublinear scaling, as expected.

$\log T + \log N$ additional qubits. Our scheme thus provides a compelling alternative to QEC in the NISQ and pre-FTQEC eras.

Hardware efficiency.—In this section, we discuss hardware overhead. First, to obtain a $1/2$ suppression, we must append a memory register and an ancilla qubit. The size of the memory register depends on the size of the *input* to $\mathcal{U}$. We stress that this is not the same as the size of $\mathcal{U}$. Besides QRAM, many useful quantum operations are possible only with an (often exponentially) large number of ancillae—fast state preparation on n qubits in $O(n)$ time requires $O(2^n/n)$ ancillae [49,50]. Since each qubit of the memory register must be swapped with the active register, this linearly increases the number of cSWAP operations. By swapping each qubit in sequentially, we can avoid adding to qubit overhead. To go from a $1/2$ to a $1/T$ suppression, one appends $\log T$ qubits to the control register only—additional error suppression has a hardware overhead that scales largely independently of $\mathcal{U}$.

A further hardware benefit is that error suppression from gate-based EF is largely agnostic to $\mathcal{U}$. This sidesteps any complications of having to construct *logical* versions of $\mathcal{U}$, since the *physical* implementation will do. In particular, error suppression via gate-based EF is independent of either the complexity of operating the black box or knowledge of the ideal unitary to be carried out. QRAM aside, state preparation also requires $\Omega(\sqrt{2^n \log n})$ T gates [51]. This non-Cliffordness compounds the hardware overhead of QEC, but does not affect gate-based EF.

Discussion.—In this work, we proposed and analyzed the circuit implementation of error filtration on a qubit quantum computer in two levels. In [31], we show that gate-based EF is isomorphic to the original error filtration setup in the noiseless case. However, our analysis considers more general error models than the loss and dephasing models considered in [18,23]. We emphasize the hardware-efficiency of gate-based EF, which allows one to use a small number of qubits with resources scaling independently of the quantum operation considered to suppress errors up to a quadratic error floor. We further emphasize that in contrast to usual error mitigation schemes, gate-based EF suppresses errors in quantum gates, which extends the reach of error mitigation to problems involving state preparation and sampling. As a comparison, one class of schemes that performs error suppression for unitary operations is the extended flag gadget scheme [52,53]. However, such schemes have limited capability in suppressing errors in non-Clifford circuits. Conversely, our scheme is entirely agnostic to the structure of the desired quantum process—in particular, QRAM is non-Clifford. Gate-based EF extends the reach of low-overhead error suppression methods, which is an important step toward bridging the NISQ and fault-tolerance eras.

While we have characterized gate-based EF by reference to NISQ-era error mitigation schemes, we note that many-controlled SWAP gates may not be easy to implement on NISQ devices. As such, one might regard gate-based EF as a scheme most suited for a post-NISQ era but before achieving FTQEC. However, this does not completely rule out the application of gate-based EF during the present NISQ era in the $T = 2$ case. In the microwave regime, a high fidelity (> 0.95) cSWAP has recently been reported in [33]. Alternatively, optical implementations may also be suitable for gate-based EF (see, e.g., [34,54,55] as well as Supplemental Material [31] for further discussion).

We thank Senrui Chen, Ng Hui Khoon, Qian Xu, and Kaiwen Gui for helpful discussions. We acknowledge support from the ARO (W911NF-18-1-0020, W911NF-18-1-0212), ARO MURI (W911NF-16-1-0349, W911NF-21-1-0325), AFOSR MURI (FA9550-19-1-0399, FA9550-21-1-0209), AFRL (FA8649-21-P-0781), DoE Q-NEXT, NSF (OMA-1936118, ERC-1941583, OMA-2137642), NTT Research, and the Packard Foundation (2020-71479).

[1] A. R. Calderbank and P. W. Shor, Good quantum error-correcting codes exist, Phys. Rev. A **54**, 1098 (1996).
[2] D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error, in Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing, *STOC'97* (Association for Computing Machinery, New York, 1997), pp. 176–188.
[3] D. Gottesman, An introduction to quantum error correction and fault-tolerant quantum computation, arXiv:0904.2557.
[4] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).
[5] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, arXiv:2210.00921.
[6] Y. Li and S. C. Benjamin, Efficient variational quantum simulator incorporating active error minimization, Phys. Rev. X **7**, 021050 (2017).
[7] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Phys. Rev. Lett. **119**, 180509 (2017), publisher: American Physical Society.
[8] C. Piveteau, D. Sutter, S. Bravyi, J. M. Gambetta, and K. Temme, Error mitigation for universal gates on encoded qubits, Phys. Rev. Lett. **127**, 200505 (2021).
[9] M. Lostaglio and A. Ciani, Error mitigation and quantum-assisted simulation in the error corrected regime, Phys. Rev. Lett. **127**, 200506 (2021).
[10] Y. Suzuki, S. Endo, K. Fujii, and Y. Tokunaga, Quantum error mitigation as a universal error reduction technique: Applications from the NISQ to the fault-tolerant quantum computing eras, PRX Quantum **3**, 010345 (2022).
[11] E. v. d. Berg, Z. K. Minev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors, arXiv:2201.09866.
[12] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, Learning-based quantum error mitigation, PRX Quantum **2**, 040330 (2021).

[13] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, Error mitigation with Clifford quantum-circuit data, Quantum **5**, 592 (2021).

[14] A. Zlokapa and A. Gheorghiu, A deep learning model for noise prediction on near-term quantum devices, arXiv:2005.10811.

[15] B. Koczor, Exponential error suppression for near-term quantum devices, Phys. Rev. X **11**, 031057 (2021).

[16] W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, Virtual distillation for quantum error mitigation, Phys. Rev. X **11**, 041036 (2021).

[17] P. Czarnik, A. Arrasmith, L. Cincio, and P. J. Coles, Qubit-efficient exponential suppression of errors, arXiv:2102.06056.

[18] N. Gisin, N. Linden, S. Massar, and S. Popescu, Error filtration and entanglement purification for quantum communication, Phys. Rev. A **72**, 012338 (2005).

[19] L.-P. Lamoureux, E. Brainis, N. J. Cerf, P. Emplit, M. Haelterman, and S. Massar, Experimental error filtration for quantum communication over highly noisy channels, Phys. Rev. Lett. **94**, 230501 (2005).

[20] G. Chiribella and H. Kristjánsson, Quantum Shannon theory with superpositions of trajectories, Proc. R. Soc. A **475**, 20180903 (2019).

[21] G. Chiribella, M. Banik, S. S. Bhattacharya, T. Guha, M. Alimuddin, A. Roy, S. Saha, S. Agrawal, and G. Kar, Indefinite causal order enables perfect quantum communication with zero capacity channels, New J. Phys. **23**, 033039 (2021).

[22] H. Kristjánsson, W. Mao, and G. Chiribella, Witnessing latent time correlations with a single quantum particle, Phys. Rev. Res. **3**, 043147 (2021).

[23] M. K. Vijayan, A. P. Lund, and P. P. Rohde, A robust *W*-state encoding for linear quantum optics, Quantum **4**, 303 (2020).

[24] Similar to our Letter, Ref. [23] proposed encoding single-qubit unitaries in a Fourier basis to similarly suppress errors. However, in our work, we show that in the gate-based context no such encoding is required, and that the resultant error suppression extends to any number of qubits.

[25] J. Guillaud and M. Mirrahimi, Repetition cat qubits for fault-tolerant quantum computation, Phys. Rev. X **9**, 041053 (2019).

[26] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais *et al.*, Bias-preserving gates with stabilized cat qubits, Sci. Adv. **6**, eaay5901 (2020).

[27] C. Chamberland, K. Noh, P. Arrangoiz-Arriola, E. T. Campbell, C. T. Hann, J. Iverson, H. Putterman, T. C. Bohdanowicz, S. T. Flammia, A. Keller, G. Refael, J. Preskill, L. Jiang, A. H. Safavi-Naeini, O. Painter, and F. G. Brandão, Building a fault-tolerant quantum computer using concatenated cat codes, PRX Quantum **3**, 010329 (2022).

[28] A. S. Darmawan, B. J. Brown, A. L. Grimsmo, D. K. Tuckett, and S. Puri, Practical quantum error correction with the XZZX code and Kerr-cat qubits, PRX Quantum **2**, 030345 (2021).

[29] D. K. Tuckett, S. D. Bartlett, and S. T. Flammia, Ultrahigh error threshold for surface codes with biased noise, Phys. Rev. Lett. **120**, 050505 (2018).

[30] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2019).

[31] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.131.190601 for derivation of analytical results, as well as discussion of the limits of gate-based EF, which includes Refs. [18,20,23,32–38].

[32] C. T. Hann, G. Lee, S. M. Girvin, and L. Jiang, Resilience of quantum random access memory to generic noise, PRX Quantum **2**, 020311 (2021).

[33] B. J. Chapman, S. J. de Graaf, S. H. Xue, Y. Zhang, J. Teoh, J. C. Curtis, T. Tsunoda, A. Eickbusch, A. P. Read, A. Koottandavida, S. O. Mundhada, L. Frunzio, M. Devoret, S. Girvin, and R. J. Schoelkopf, High-on-off-ratio beamsplitter interaction for gates on bosonically encoded qubits, PRX Quantum **4**, 020355 (2023).

[34] K. C. Chen, W. Dai, C. Errando-Herranz, S. Lloyd, and D. Englund, Scalable and high-fidelity quantum random access memory in spin-photon networks, PRX Quantum **2**, 030319 (2021).

[35] R. Chao and B. W. Reichardt, Quantum error correction with only two extra qubits, Phys. Rev. Lett. **121**, 050502 (2018).

[36] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum random access memory, Phys. Rev. Lett. **100**, 160501 (2008).

[37] F.-Y. Hong, Y. Xiang, Z.-Y. Zhu, L.-z. Jiang, and L.-n. Wu, Robust quantum random access memory, Phys. Rev. A **86**, 010306(R) (2012).

[38] L.-M. Duan and H. J. Kimble, Scalable photonic quantum computation through cavity-assisted interactions, Phys. Rev. Lett. **92**, 127902 (2004).

[39] U. Réglade, A. Bocquet, R. Gautier, A. Marquet, E. Albertinale, N. Pankratova, M. Hallén, F. Rautschke, L.-A. Sellem, P. Rouchon, A. Sarlette, M. Mirrahimi, P. Campagne-Ibarcq, R. Lescanne, S. Jezouin, and Z. Leghtas, Quantum control of a cat-qubit with bit-flip times exceeding ten seconds, arXiv:2307.06617.

[40] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret, Stabilization and operation of a Kerr-cat qubit, Nature (London) **584**, 205 (2020).

[41] Q. Xu, G. Zheng, Y.-X. Wang, P. Zoller, A. A. Clerk, and L. Jiang, Autonomous quantum error correction and fault-tolerant quantum computation with squeezed cat qubits, npj Quantum Inf. **9**, 78 (2023).

[42] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, Nature (London) **508**, 500 (2014).

[43] Note that log is taken to base 2 in this text, such that $\log T$ refers to the number of control qubits utilized.

[44] V. Giovannetti, S. Lloyd, and L. Maccone, Architectures for a quantum random access memory, Phys. Rev. A **78**, 052310 (2008).

[45] S. Arunachalam, V. Gheorghiu, T. Jochym-O'Connor, M. Mosca, and P. V. Srinivasan, On the robustness of bucket brigade quantum ram, New J. Phys. **17**, 123010 (2015).

[46] T. J. Yoder, R. Takagi, and I. L. Chuang, Universal fault-tolerant gates on concatenated stabilizer codes, Phys. Rev. X **6**, 031039 (2016).

[47] C. D. Hill, A. G. Fowler, D. S. Wang, and L. C. L. Hollenberg, Fault-tolerant quantum error correction code conversion, arXiv:1112.2417.

[48] O. D. Matteo, V. Gheorghiu, and M. Mosca, Fault-tolerant resource estimation of quantum random-access memories, IEEE Trans. Quantum Eng. **1**, 1 (2020).

[49] K. Gui, A. M. Dalzell, A. Achille, M. Suchara, and F. T. Chong, Spacetime-efficient low-depth quantum state preparation with applications, arXiv:2303.02131.

[50] P. Yuan and S. Zhang, Optimal (controlled) quantum state preparation and improved unitary synthesis by quantum circuits with any number of ancillary qubits, Quantum **7**, 956 (2023).

[51] G. H. Low, V. Kliuchnikov, and L. Schaeffer, Trading *t*-gates for dirty qubits in state preparation and unitary synthesis, arXiv:1812.00954.

[52] D. M. Debroy and K. R. Brown, Extended flag gadgets for low-overhead circuit verification, Phys. Rev. A **102**, 052409 (2020).

[53] A. Gonzales, R. Shaydulin, Z. H. Saleem, and M. Suchara, Quantum error mitigation by Pauli check sandwiching, Sci. Rep. **13**, 2122 (2023).

[54] F.-Y. Hong, Y. Xiang, Z.-Y. Zhu, L.-z. Jiang, and L.-n. Wu, Robust quantum random access memory, Phys. Rev. A **86**, 010306 (2012).

[55] T. Ono, R. Okamoto, M. Tanida, H. F. Hofmann, and S. Takeuchi, Implementation of a quantum controlled-SWAP gate with photonic circuits, Sci. Rep. **7**, 45353 (2017).