

ARTICLE OPEN

Accounting for errors in quantum algorithms via individual error reduction

Matthew Otten¹ and Stephen K. Gray¹

We discuss a surprisingly simple scheme for accounting (and removal) of error in observables determined from quantum algorithms. A correction to the value of the observable is calculated by first measuring the observable with all error sources active and subsequently measuring the observable with each error source reduced separately. We apply this scheme to the variational quantum eigensolver, simulating the calculation of the ground state energy of equilibrium H₂ and LiH in the presence of several noise sources, including amplitude damping, dephasing, thermal noise, and correlated noise. We show that this scheme provides a decrease in the needed quality of the qubits by up to two orders of magnitude. In near-term quantum computing, where full fault-tolerant error correction is too expensive, this scheme provides a route to significantly more accurate calculations.

npj Quantum Information (2019)5:11; <https://doi.org/10.1038/s41534-019-0125-3>

INTRODUCTION

Quantum computing is beginning to show promising proof-of-principle calculations, especially in quantum chemistry. Calculations of the binding energies for molecules such as H₂¹ and BeH₂² have been done using small, noisy quantum computers. Applications in machine learning have also been shown on quantum hardware.³ Quantum computing is entering the noisy intermediate-scale quantum era.⁴ Full fault-tolerant error correction is still many years away; near-term quantum computers will have a limited number of qubits, each qubit being noisy. Methods that reduce noise and correct errors without doing full error correction on every qubit will help extend the range of interesting problems that can be solved in the near-term. Each qubit is a very valuable resource for near-term quantum computers; using them efficiently, and reusing them as often as possible, is imperative in enabling useful calculations on near-term devices.

Here we describe and demonstrate a simple scheme for reducing the effects of a variety of noise sources by reducing each source separately and summing the resulting corrections. In practice, the reduction could be accomplished via quantum error correction⁵ or by active engineering to reduce a noise rate, e.g., autonomous error correction^{6,7} or dynamical decoupling.⁸ Simple quantum error correction schemes have already been shown in superconducting circuits^{9,10} and such systems are approaching more complicated error correction schemes.¹¹ In trapped ion quantum computers, a seven qubit color code has been already been demonstrated.¹² Quantum computing architectures are nearing the quality and size where a single qubit could be error corrected, but we are far from every qubit being corrected. The scheme we present could make use of this limited error correction, by sweeping through each qubit, correcting one qubit at a time, but is not dependent on error correction methods per se. We validate it with the variational quantum eigensolver (VQE),^{13,14} simulating the calculation of the ground state energies of H₂ and LiH. We assume that a single qubit has its error rate reduced by some means, while the other qubits retain all of their error. Our

results show that the scheme reduces the needed quality of each qubit drastically; for ‘chemical accuracy’, error rates can be up to two orders of magnitude larger. We apply this scheme to multiple noise sources, including amplitude damping, dephasing, thermal noise, and correlated noise. We stress that the scheme can be used to reduce the environmental error from any measured observable, not just those used in VQE. Furthermore, this scheme is not restricted to algorithms on quantum computers; it could potentially be used in quantum sensing. Quantum error correction has been proposed as a method to reduce environmental noise in quantum metrology,^{7,15} allowing a probe to reach the Heisenberg limit.¹⁶ Our scheme could reduce the number of error correcting qubits necessary for such a method.

RESULTS

Noise sources and removal

We simulate four different types of varying noise sources, represented by Lindblad operators: amplitude damping (L^1), dephasing (L^2), thermal (L^{th}), and correlated noise (L^c):

$$\begin{aligned} L^1(\rho) &= \gamma_1 \mathcal{D}[\sigma](\rho), \\ L^2(\rho) &= \gamma_2 \mathcal{D}[\sigma^\dagger \sigma](\rho), \\ L^{th}(\rho) &= \gamma_{th} (n_{th} + 1) \mathcal{D}[\sigma] + \gamma_{th} n_{th} \mathcal{D}[\sigma^\dagger], \\ L^c(\rho) &= \gamma_c \mathcal{D}[\sigma_1^\dagger \sigma_2] + \gamma_c \mathcal{D}[\sigma_1 \sigma_2^\dagger], \end{aligned} \quad (1)$$

where σ is an annihilation operator and $\mathcal{D}[C](\rho) = C\rho C^\dagger - \frac{1}{2}(C^\dagger C\rho + \rho C^\dagger C)$. These Lindblad terms are applied to each qubit or to various combinations of qubits. The parameters in Eq. (1) are $\gamma_1 = \frac{1}{T_1}$, the amplitude damping rate; $\gamma_2 = \frac{2}{T_2}$, twice the dephasing rate; γ_{th} , the thermalization rate; n_{th} , the thermal occupation (taken to be 0.5); and γ_c , the correlated noise rate. We have used this formalism in previous work.^{17–19} For our calculations we assume that reduction of an error source corresponds to reducing the relevant parameter, γ , by some known amount.

¹Center for Nanoscale Materials, Argonne National Laboratory, Lemont, IL 60439, USA
Correspondence: Matthew Otten (otten@anl.gov)

Received: 5 July 2018 Accepted: 3 January 2019

Published online: 29 January 2019

It is important to understand the exact process of the reduction of the relevant error parameters and what effect that has on the relevant error properties of the underlying qubits. For instance, if some active engineering process reduces the error only fractionally, the fractional reduction must be known to make full use of our method; if it is only known approximately, there will still be a correction, but it will not be as drastic. Furthermore, if the error reduction strategy decreases the relevant source of error, but increases a different source of error, the method will not directly be able to account for the additional error. For instance, quantum error correction could be used to reduce the error by correcting each qubit, one at a time. Quantum error correction, however, naturally introduces new overhead in number of qubits (due to the encoding) and in number of gates (due to the increased complexity of logical operations, especially two-qubit operations, on encoded qubits), which could introduce new types of errors into the circuit evaluation. This additional complexity would make the original, noisy evaluation differ by more than just a reduction of the error on each qubit by a partial amount. Furthermore, the quantum error correction algorithms likely to be implemented in the near-term will only partially decrease the error, not eliminate it, even without all these complications. In this work, we assume gates are separated by one time unit and the error rates are given in inverse time units.

Variational quantum eigensolver

Here, we briefly overview the variational quantum eigensolver (VQE). VQE solves for an approximate, variational ground state energy of a parameterized wavefunction ansatz, $|\psi(\theta)\rangle$. The variational principle ensures that E_0 , the true ground state energy of the Hamiltonian H , is always less than E , the energy of a parameterized wavefunction ansatz. E is evaluated on the quantum computer; the parameters θ are optimized using a classical computer. Classical computing methods such as variational quantum Monte Carlo²⁰ also make use of the variational principle. The hope of a quantum realization is that quantum computers can efficiently prepare non-trivial states which would be more difficult to prepare on a classical computer. While methods like quantum phase estimation²¹ can give generally more accurate energies, VQE requires shorter circuits and has a natural robustness to noise.^{2,13,22} VQE is not limited to quantum chemistry; it has also been used to study problems in nuclear physics.²³ When using VQE for quantum chemistry, the second quantized quantum chemistry Hamiltonian is transformed into a qubit Hamiltonian using a transformation such as Jordan–Wigner.²⁴

Correction scheme

First consider the Lindblad master equation,

$$\frac{d\rho}{dt} = L(\rho) = \sum_{i=1}^m L_i(\rho). \quad (2)$$

A formal solution for a density matrix evolving from time t to $t + \tau$ and satisfying this equation is

$$\rho(t + \tau) = V_\tau(\rho(t)), \quad (3)$$

where

$$V_\tau(\cdot) = \exp[\tau L(\cdot)]. \quad (4)$$

We use $()$ above to indicate that L and V_τ are superoperators that take in an operator with the brackets to generate a new one. To first order in τ and taking $L()$ to be the sum over Lindblad operators of Eq. (2),

$$V_\tau(\cdot) \approx 1 + \tau \sum_{i=1}^m L_i(\cdot), \quad (5)$$

Applying gate 1 via application of unitary operator U_1 , evolving under V_τ , applying gate 2 (U_2), etc., up to gate G (U_G) corresponds exactly to a final density matrix given by

$$\rho(T) = U_G V_\tau (U_{G-1} \cdots V_\tau (U_2 V_\tau (U_1 \rho(0) U_1^\dagger) U_2^\dagger) \cdots U_{G-1}^\dagger) U_G^\dagger, \quad (6)$$

where $T = (G - 1) \tau$, the time after the last gate has been applied. Notice that Eq. (6) is not symmetric, since V_τ is a superoperator.

Equation (6) is $\rho(T)$ for the case of all m error sources present. The corresponding $\rho_a(T)$ with no error terms present is simply:

$$\rho_a(T) = U_G U_{G-1} \cdots U_2 U_1 \rho(0) U_1^\dagger U_2^\dagger \cdots U_{G-1}^\dagger U_G^\dagger. \quad (7)$$

The density matrices resulting from reducing error sources $i = 1, 2, \dots, m$ separately are

$$\rho_i(T) = U_G V_\tau^i (U_{G-1} \cdots V_\tau^i (U_2 V_\tau^i (U_1 \rho(0) U_1^\dagger) U_2^\dagger) \cdots U_{G-1}^\dagger) U_G^\dagger. \quad (8)$$

V_τ^i is the corresponding Lindblad evolution operator that contains a suitably scaled L_i but has all other terms the same, e.g. to first order in τ ,

$$V_\tau^i(\cdot) \approx 1 + \tau \sum_{j \neq i}^m L_j(\cdot) + \tau(1 - f_i) L_i(\cdot), \quad (9)$$

where f_i is the fraction of noise removed.

Now we consider the new density matrix defined as

$$\tilde{\rho}(T) = \rho(T) - \sum_{i=1}^m \frac{1}{f_i} (\rho(T) - \rho_i(T)). \quad (10)$$

Insertion of the Eqs. (5) and (9) into Eq. (10) leads, as detailed in the Supplementary Material, to

$$\tilde{\rho}(T) = \rho_a(T) + \mathcal{O}(\tau^2) \quad (11)$$

being correct to first order in τ , i.e. all first order error terms exactly cancel, with remaining error terms on the order of τ^2 and higher. In contrast, the uncorrected density matrix $\rho(T)$ contains first order error terms.

We should note that Eq. (10) could also be obtained via a linear expansion of the density matrix about the noise rates and a finite difference approximation to the first derivatives. However, we prefer our propagator-based derivation since it provides a more physical context as well as a means for rigorous error analysis in terms of the underlying propagation and Lindblad noise terms. The general philosophy of error mitigation via additional measurements and linear response can also be found in somewhat more sophisticated approaches such as quantum subspace expansion.²⁵

Since the density matrix itself is corrected by this procedure, all observables are also accurate to first order in τ . Equation (10) immediately leads to the observable correction formula, traced with the observable, A :

$$\begin{aligned} \tilde{A} &= \text{Tr}(\tilde{\rho}(T)A) \\ &= \text{Tr}(\rho(T)A) - \sum_{i=1}^m \frac{1}{f_i} [\text{Tr}(\rho(T)A) - \text{Tr}(\rho_i(T)A)] \\ &= \langle A \rangle - \sum_{i=1}^m \frac{1}{f_i} (\langle A \rangle - \langle A_i \rangle), \end{aligned} \quad (12)$$

which will also be accurate to order τ whereas the uncorrected observable, $\langle A \rangle$, has first order error terms.

Suppose, for example, that $m = n$, where n is the number of qubits, and each qubit is noisy. The strength of Eq. (12) is that, for near-term quantum computing without the possibility of perfect error correction of all n qubits, only $\mathcal{O}(n)$ computations involving reducing the error on just one qubit (to yield the $\langle A_i \rangle$) are required, along with the original calculation with no error reduction to yield $\langle A \rangle$. Intuitively, the difference $f_i(\langle A \rangle - \langle A_i \rangle)$ represents a fraction of the first order noise contributions for noise source i ; these are then subtracted away from the noisy expectation value with a relevant

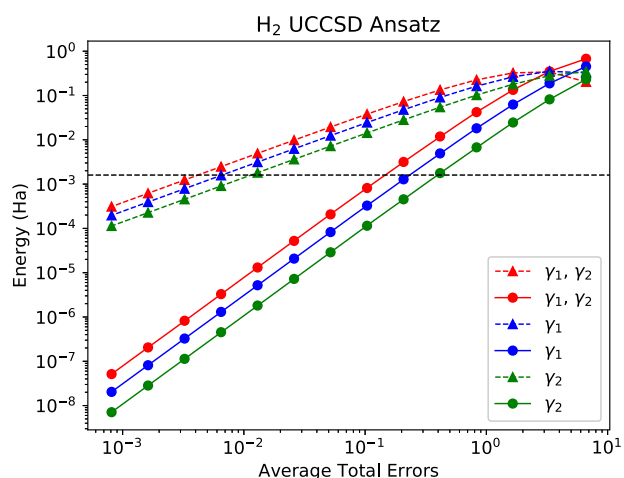


Fig. 1 Error cancellation for ground state H_2 with the UCCSD ansatz under amplitude damping (γ_1), dephasing (γ_2), and both amplitude damping and dephasing (γ_1, γ_2) noise sources, removing only 10% of the noise on each qubit separately. The horizontal line represents 'chemical accuracy', 1.6 mHa. Dashed lines with triangles are the amount of correction applied by our scheme. Solid lines with circles are the difference between the corrected energy and the energy evaluated with no noise

scaling $\frac{1}{f_i}$, leaving a result with less noise. This cancellation relies on the expectation value being built up from many measurements. Each expectation value ($\langle A \rangle$ and $\langle A_i \rangle$) contains contributions from measurements with no error, as well as measurements with errors. Our correction scheme cancels out some of the measurements with error, while leaving the result with no error, leading to a better calculation of the observable. Our simulations of the application of this scheme to VQE indicate that Eq. (12) does indeed yield a substantial improvement for real algorithms. Our method lends itself naturally to use on a quantum computer, where calculations can be repeated and techniques exist for the reduction of error, but it is not restricted to just that. Any quantum system in which an observable can be repeatedly measured and each noise source can be reduced separately can make use of the scheme to obtain a more accurate result; quantum error correction assisted quantum metrology is a prime example.^{7,15}

In VQE, the measured observable in question is the energy, E , of the wavefunction ansatz. We first optimize the parameters of the wavefunction; this can be done either with no error reduction or potentially reducing the error on a single qubit. Once a set of optimal parameters is found, the expectation value of the energy is evaluated on the quantum computer with no error reduction, and then reducing the error on each qubit separately. For an n qubit problem, this involves only an additional $\mathcal{O}(n)$ evaluations of the energy on the quantum computer with error reduction on one of the qubits each time. Once all of the energies are measured, Eq. (12) is used to obtain \bar{E} .

H_2

We first consider the hydrogen molecule, H_2 , at equilibrium bond length 0.74 Å and an STO-3G basis, resulting in a four qubit circuit. We use the unitary coupled cluster singles double (UCCSD) ansatz¹ (166 gates) and note that each gate is applied sequentially with one time unit between each gate. We made no effort to apply gates in parallel. The parameters of the wavefunction ansatz were optimized with noise on every qubit. We then sweep through the qubits, reducing the noise from each qubit by 10% ($f_i = 0.1$ for each noise source i). The final energy is then calculated by using our correction scheme, Eq. (12).

Results for typical amplitude damping (γ_1) and dephasing (γ_2) noises are shown in Fig. 1, representing three different environmental regimes. One regime corresponds to $\gamma_1 = \gamma_2$ and is similar to a superconducting qubit quantum computer,² whereas the regime with just γ_2 is consistent with spin²⁶ and trapped ion²⁷ quantum computers. Though we do not know of a quantum architecture where γ_1 is the dominant noise source, we include it as a third regime for completeness. The x-axis in this figure (and all subsequent figures) shows the error rate multiplied by the number of qubits, and roughly represents the expected number of errors for a given circuit evaluation. For the $\gamma_1 = \gamma_2$ case, however, the expected number of errors would be doubled, as there are two possible error channels. We see that chemical accuracy (1.6 mHa, the horizontal black line) with respect to the noise-free result can be obtained with error rates more than one order of magnitude higher than without our method; on average, the error rates can be 35× larger. To obtain chemical accuracy, corrections of 50 mHa are applied. On this log-log plot, the slopes of the two lines represent their scaling with error. The correction amount (which is approximately the error of the uncorrected energy) has a slope of 1, whereas the corrected energy has a slope of 2, representing the cancellation of the first order error terms.

The Supplementary Material provides results for this same example removing 100% of the noise. This results in only a slight benefit, compared to removing only a small fraction (10%) of the noise, allowing for error rates to be, on average, 45× larger (30% larger than the fractional case). By removing all of the noise, there is no need for the linear extrapolation done by the factor $\frac{1}{f_i}$, the value at zero-noise is directly measured for each qubit. Additionally, the determination of the fraction of removed noise, f , can be complicated. This could be done by measuring the relevant error properties before and after the error reduction strategy is applied, but errors in the process could give an unreliable estimation of f . Even a small deviation from the true value of f will eventually result in the slope of the corrected energy becoming 1 (instead of 2) at sufficiently small error rates. This is due to the difference in the corrected energy and the noise-free energy becoming smaller than the precision of the determined f . If an unknown f is assumed to be 1, for instance, the method will still remove a fraction of the error, giving results as if the quantum computer had error rates $1-f$; there are still first order errors, but they are smaller due to the correction scheme. Because the difference between the result when a small fraction of the noise is removed and when all of the noise is removed is relatively marginal, we focus on the latter case for the remainder of this paper, where Eq. (12) takes a simpler form with $f=1$. The Supplementary Material also provides results for a different wavefunction ansatz, one similar to that of ref.². The results are consistent when using this separate ansatz. Note that these figures make no reference to the true, full configuration interaction energy. The difference plotted is between the corrected energy and the energy evaluated with no noise; the quality of the wavefunction ansatz cannot be determined from these figures alone. The ordering of the results provides a limited sensitivity analysis for different quantum computing architectures. Similar to ref.²², we note that VQE is more sensitive to amplitude damping noise than to dephasing noise.

We should note that the very largest gains in accuracy in Fig. 1 (and subsequent figures) occur at relatively low error rates. The initial, uncorrected algorithm that is run should be reasonably reliable, e.g., the error rate should be $1/(\text{no. gates})$, in inverse time units with gates separated by one time unit. This error rate (e.g., $1/166 \approx 0.006$ for the case of Fig. 1) may be quite daunting to achieve in practice; partial error correction on all qubits could be used to achieve this threshold.

Even though the wavefunction parameters were optimized in the presence of noise, the final energy evaluated at the different

parameter sets for the fully error corrected circuit differ very little. The optimal parameters from the largest error rates only give a difference of 1.7 mHa compared to the optimal parameters from the error-free optimal parameters, when both were evaluated with no noise. We therefore optimize the parameters once with no noise and use those parameters for evaluation at all noise rates in the following examples.

Our correction scheme is not limited to environmental noise sources, such as those modeled by γ_1 and γ_2 . Any noise source describable by a Lindblad superoperator can benefit, as long as each noise source can be isolated and reduced independently of all other noise sources. To demonstrate this, we apply a thermal noise source with rate γ_{th} and a correlated noise source with rate γ_c to the H_2 UCCSD example, Fig. 2, where we now assume that the noise has been completely removed ($f_i = 1$ for all i). We see trends similar to those for amplitude damping and dephasing; corrections of ≈ 70 mHa bring the energy to within chemical accuracy at error rates almost 50 times larger than otherwise needed. While perhaps experimentally difficult, thermal noise could be reduced by selectively cooling each qubit. The correction

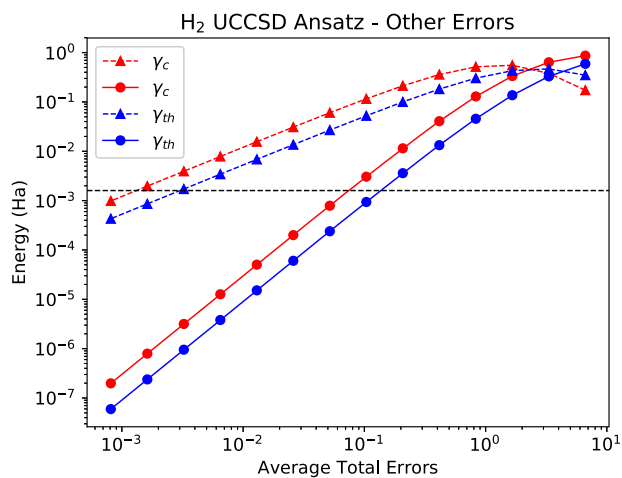


Fig. 2 Error cancellation for ground state H_2 with the UCCSD ansatz under thermal noise (γ_{th}) and correlated noise (γ_c) with total error removal on each qubit separately. See caption of Fig. 1 for explanation of linetypes and symbols

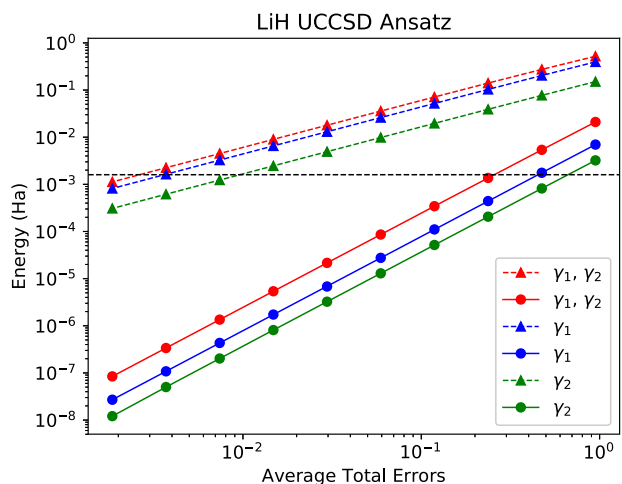


Fig. 3 Error cancellation for the ground state of LiH with the UCCSD ansatz under amplitude damping (γ_1), dephasing (γ_2), and both amplitude damping and dephasing (γ_1, γ_2) noise sources with total error removal on each qubit separately. See caption of Fig. 1 for explanation of linetypes and symbols

scheme applied to our correlated noise term reveals some subtleties of the method. Our correlated noise Lindblad, L^c of Eq. (1), naturally has terms from two qubits. When we sweep through the qubits, we now reduce fully all terms which involve a single qubit; this leads to the removal of each L^c term twice, once for each qubit in each L^c . This double counting can be corrected by taking half of the calculated correction from each qubit. Our scheme relies on the fact that each term is reduced only once. As long as the noise sources of interest and their controlled reduction are well understood, the scheme can be applied. In our correlated noise term, L^c , every term is removed exactly twice and the calculated correction can be halved. Though we focus on fractional noise reduction, the results will still hold if the noise is instead increased by a controlled, known amount (say, doubled), for each qubit separately. The ‘correction’ would be the difference between the inflated noise run and the normal noise run, scaled by the appropriate factor. This is similar in spirit to refs. ^{28–30}, where the total noise of the system is artificially increased and the results are subsequently extrapolated to the zero noise limit.

LiH

We also study LiH in the STO-3G basis at bond length 1.74 Å, using 12 qubits, with over 12,000 gates. Results for the UCCSD ansatz are shown in Fig. 3. The correction for LiH is even more dramatic than for H_2 . Corrections of 100–200 mHa bring the answer to within chemical accuracy, and error rates can be over two orders of magnitude higher, ranging from 68 for only γ_1 noise to 128 with only γ_2 noise. This example provides confidence that the correction scheme will work for larger circuits. For LiH the procedure works even better than for the smaller circuits of H_2 . This will likely be true for ever larger circuits: the number of first order errors increases with increasing number of qubits, and these are all approximately corrected. While the number of second order errors also increases, second order errors go as γ^2 and τ^2 , and will be small.

Our method can reduce the error in any measured observable, and so has application to a wide range of quantum algorithms. Further study on other algorithms, e.g., the quantum phase estimation²¹ and quantum approximate optimization algorithms,³¹ are needed to understand the broader impact of this scheme. Applications to quantum metrology warrant further study: the noise floor of a quantum sensor could be reduced at substantial reduction in cost in number of quantum systems by reducing the overhead for quantum error correction assisted metrology.^{7,15} The magnitude of the correction can also be used as a metric for measuring how close to the true answer one is, without knowledge of the true answer. The method relies on the error characteristics with no reduction to be the same as the error characteristics after reduction; in quantum error correction, this may not always be true. Further study into practical methods of reducing the error and ensuring that the error characteristics before and after are similar is ongoing.

DISCUSSION

We presented a simple scheme to reduce error in quantum algorithms, applying it to simulations of the variational quantum eigensolver. Reducing the error on each qubit (e.g., through error correction or some active or engineering process) one at a time and summing the scaled difference from the result with no error reduction can provide a significant correction to the observable (e.g., energy). This scheme reduces the coherence requirements to obtain chemical accuracy; error rates can be up to two orders of magnitude larger. The overhead is relatively low: an additional $\mathcal{O}(n)$ evaluations with single qubit error reduction.

The general philosophy used in this work for accounting for errors, replacing all-qubit error correction with additional

measurements and making use of an error model, has also been exploited in interesting recent work^{28–30} involving increasing a single global noise parameter, along with extrapolation to zero error assuming a polynomial expansion for the error or via a quasi-probability formalism. Our approach focuses on changing the error associated with each individual qubit or error source separately. Furthermore, our method is developed with focus on mitigating the environmental noise of each qubit, whereas these other interesting works focus on the noise from gate applications. In the near term, with a very limited number of qubits and gates, it is likely that the approach of refs^{28–30} is more feasible. However, our method, presented in the paper, could be used for quantum memories and other quantum devices, where the dominant noise source is not the gates, but the environmental interactions. Our approach offers an alternative that is particularly relevant when increasing a global noise parameter (such as the noise on the gates) is infeasible. Other, more recent work, focuses on using a stabilizer formalism, where a specific, relevant quantity (such as a global symmetry of the system) is used to detect errors occurring in a quantum circuit.^{32,33} If the symmetry is violated, then an error occurred and the run in question should be excluded. Such approaches cannot eliminate all errors and could be used in tandem with the method we propose in this paper. As quantum devices with larger numbers of qubits come online, utilization of many forms of error mitigation will be necessary to make full use of the power of these small, noisy devices.

METHODS

Generation of VQE circuits

We use the open source package OpenFermion³⁴ to generate the qubit Hamiltonian, starting from quantum chemistry integrals generated via Psi4.³⁵ We use the unitary coupled cluster singles doubles (UCCSD) ansatz,¹³ with OpenFermion³⁴ and ProjectQ³⁶ to generate the circuits. We optimize the parameters of the wavefunctions using both Nelder-Mead simplex³⁷ and COBYLA.³⁸

Time evolution

Consider a system of n qubits characterized by a time-dependent density matrix $\rho(t)$ subjected to a sequence of $k = 1, 2, \dots, G$ gate operations, each being a unitary transformation U_k on ρ : $\rho \rightarrow U_k \rho U_k^\dagger$. We assume time τ lapses between each gate operation. During these times ρ evolves under a Lindblad master equation, Eq. (2). The dynamics is simulated with the high-performance density matrix evolution program Quac,³⁹ using the different noise sources noted in the main text.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author on reasonable request.

ACKNOWLEDGEMENTS

This work was performed at the Center for Nanoscale Materials, a U.S. Department of Energy Office of Science User Facility, and supported by the U.S. Department of Energy, Office of Science, under Contract No DE-AC02-06CH11357. We gratefully acknowledge the computing resources provided on Bebop, a high-performance computing cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory.

AUTHOR CONTRIBUTIONS

M.O. and S.K.G. contributed to the research and wrote the paper. M.O. designed the method and performed the numerical simulations.

ADDITIONAL INFORMATION

Supplementary information accompanies the paper on the *npj Quantum Information* website (<https://doi.org/10.1038/s41534-019-0125-3>).

Competing interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

REFERENCES

- O'Malley, P. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
- Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242 (2017).
- Otterbach, J. et al. Unsupervised machine learning on a hybrid quantum computer. arXiv:1712.05771 (2017).
- Preskill, J. Quantum computing in the nisq era and beyond. arXiv:1801.00862 (2018).
- Nielsen, M. A. & Chuang, I. L. *Quantum computation and quantum information* (Cambridge University Press, Cambridge, UK, 2010).
- Freeman, C. D., Herdman, C. M. & Whaley, K. B. Engineering autonomous error correction in stabilizer codes at finite temperature. *Phys. Rev. A* **96**, 012311 (2017).
- Reiter, F., Sørensen, A. S., Zoller, P. & Muschik, C. Dissipative quantum error correction and application to quantum sensing with trapped ions. *Nat. Commun.* **8**, 1822 (2017).
- Souza, A. M., Álvarez, G. A. & Suter, D. Robust dynamical decoupling for quantum computing and quantum memory. *Phys. Rev. Lett.* **106**, 240501 (2011).
- Reed, M. D. et al. Realization of three-qubit quantum error correction with superconducting circuits. *Nature* **482**, 382 (2012).
- Kelly, J. et al. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature* **519**, 66 (2015).
- Barends, R. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **508**, 500 (2014).
- Nigg, D. et al. Quantum computations on a topologically encoded qubit. *Science* **345**, 302–305 (2014).
- Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).
- McClean, J. R., Romero, J., Babbush, R. & Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **18**, 023023 (2016).
- Kessler, E. M., Lovchinsky, I., Sushkov, A. O. & Lukin, M. D. Quantum error correction for metrology. *Phys. Rev. Lett.* **112**, 150802 (2014).
- Wineland, D. J., Bollinger, J. J., Itano, W. M., Moore, F. L. & Heinzen, D. J. Spin squeezing and reduced quantum noise in spectroscopy. *Phys. Rev. A* **46**, R6797–R6800 (1992).
- MacQuarrie, E. R., Otten, M., Gray, S. K. & Fuchs, G. D. Cooling a mechanical resonator with nitrogen-vacancy centres using a room temperature excited state spin-strain interaction. *Nat. Commun.* **8**, 14358 (2017).
- Otten, M. et al. Origins and optimization of entanglement in plasmonically coupled quantum dots. *Phys. Rev. A* **94**, 022312 (2016).
- Otten, M. et al. Entanglement of two, three, or four plasmonically coupled quantum dots. *Phys. Rev. B* **92**, 125432 (2015).
- Nightingale, M. P. & Umrigar, C. J. *Quantum Monte Carlo methods in physics and chemistry*. 525 (Springer Science & Business Media, Dordrecht, The Netherlands, 1998).
- Kitaev, A. Y. Quantum measurements and the abelian stabilizer problem. arXiv: quant-ph/9511026 (1995).
- Sawaya, N. P., Smelyanskiy, M., McClean, J. R. & Aspuru-Guzik, A. Error sensitivity to environmental noise in quantum circuits for chemical state preparation. *J. Chem. Theory Comput.* **12**, 3097–3108 (2016).
- Dumitrescu, E. F. et al. Cloud quantum computing of an atomic nucleus. *Phys. Rev. Lett.* **120**, 210501 (2018).
- Whitfield, J. D., Biamonte, J. & Aspuru-Guzik, A. Simulation of electronic structure hamiltonians using quantum computers. *Mol. Phys.* **109**, 735–750 (2011).
- McClean, J. R., Kimchi-Schwartz, M. E., Carter, J. & de Jong, W. A. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A* **95**, 042308 (2017).
- Watson, T. et al. A programmable two-qubit quantum processor in silicon. *Nature* **555**, 633–637 (2018).
- Kielinski, D., Monroe, C. & Wineland, D. J. Architecture for a large-scale ion-trap quantum computer. *Nature* **417**, 709 (2002).
- Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017).
- Temme, K., Bravyi, S. & Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, 180509 (2017).

30. Kandala, A. et al. Extending the computational reach of a noisy superconducting quantum processor. arXiv:1805.04492 (2018).
31. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. arXiv:1411.4028 (2014).
32. Bonet-Monroig, X., Sagastizabal, R., Singh, M. & O'Brien, T. Low-cost error mitigation by symmetry verification. arXiv:1807.10050 (2018).
33. McArdle, S., Yuan, X. & Benjamin, S. Error mitigated quantum computational chemistry. arXiv:1807.02467 (2018).
34. McClean, J. R. et al. Openfermion: the electronic structure package for quantum computers. arXiv:1710.07629 (2017).
35. Parrish, R. M. et al. Psi4 1.1: an open-source electronic structure program emphasizing automation, advanced libraries, and interoperability. *J. Chem. Theory Comput.* **13**, 3185–3197 (2017).
36. Steiger, D. S., Häner, T. & Troyer, M. Projectq: an open source software framework for quantum computing. *Quantum* **2**, 49 (2018).
37. Nelder, J. A. & Mead, R. A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965).
38. Powell, M. J. A *Advances in optimization and numerical analysis*, 51–67 (Springer, 1994).
39. Otten, M. Quac: open quantum systems in C, a time-dependent open quantum systems solver. <https://github.com/Ott3r/QuaC> (2017).



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

This is a U.S. government work and not under copyright protection in the U.S.; foreign copyright protection may apply 2019