Low-cost error mitigation by symmetry verification

Bonet-Monroig, X.; Sagastizabal, R.; Singh, M.; O'Brien, T. E.

DOI

[10.1103/PhysRevA.98.062339](https://doi.org/10.1103/PhysRevA.98.062339)

Publication date

2018

Document Version

Final published version

Published in

Physical Review A

Citation (APA)

Bonet-Monroig, X., Sagastizabal, R., Singh, M., & O'Brien, T. E. (2018). Low-cost error mitigation by symmetry verification. *Physical Review A*, 98(6), [062339]. <https://doi.org/10.1103/PhysRevA.98.062339>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Low-cost error mitigation by symmetry verification

X. Bonet-Monroig,¹ R. Sagastizabal,^{2,3} M. Singh,^{2,3} and T. E. O'Brien¹

¹*Instituut-Lorentz, Universiteit Leiden, P.O. Box 9506, NL-2300 RA Leiden, The Netherlands*

²*QuTech, Delft University of Technology, P.O. Box 5046, NL-2600 GA Delft, The Netherlands*

³*Kavli Institute of Nanoscience, Delft University of Technology, P.O. Box 5046, NL-2600 GA Delft, The Netherlands*

(Received 9 August 2018; published 28 December 2018)

We investigate the performance of error mitigation via measurement of conserved symmetries on near-term devices. We present two protocols to measure conserved symmetries during the bulk of an experiment, and develop a third, zero-cost, post-processing protocol which is equivalent to a variant of the quantum subspace expansion. We develop methods for inserting global and local symmetries into quantum algorithms, and for adjusting natural symmetries of the problem to boost the mitigation of errors produced by different noise channels. We demonstrate these techniques on two- and four-qubit simulations of the hydrogen molecule (using a classical density-matrix simulator), finding up to an order of magnitude reduction of the error in obtaining the ground-state dissociation curve.

DOI: [10.1103/PhysRevA.98.062339](https://doi.org/10.1103/PhysRevA.98.062339)

I. INTRODUCTION

Noisy, intermediate scale quantum (NISQ) devices have begun to appear in laboratories around the world. These devices have performance rates around or just below the quantum error correction threshold [1–5], but are lacking the number of qubits required for full fault-tolerant quantum computing. This raises the open question of whether the upcoming generation of quantum computers will provide a quantum advantage over classical computers, and in which fields this might be achieved [6–8]. In particular, for the area of digital quantum simulation, it has been suggested that variational quantum eigensolvers [9] may be sufficiently low cost to be performed on ~ 50 qubits [10–13]. Around this point, solving the many-body problem exactly becomes too challenging for classical computers, and a slight quantum edge might be available above current approximations.

In lieu of full error correction techniques, much attention is being turned to error mitigation techniques, which, although unscalable, promise modest improvements at low cost. Previous work has focused on active error minimization, whereby data is obtained at artificially increased error rates and then extrapolated to zero [14–17], and on probabilistic error cancellation, where an ensemble of noisy circuits is applied such that they average to the target error-free circuit [14,18]. More specific techniques have been developed for quantum simulation, and in particular for variational quantum eigensolvers. A technique developed for exploring the low-energy excited subspace of a quantum system, the quantum subspace expansion, has been shown to have error mitigation as a side effect [19,20].

In this work we investigate error mitigation via verification of symmetries found in quantum circuits, in particular those in physical systems. This is a low-cost version of the stabilizer parity checks ubiquitous in quantum error correction [21,22]. We develop multiple protocols to perform symmetry verification, both repeatedly throughout a quantum circuit and as

a single post-processing step. The latter can be related to a variant of the quantum subspace expansion [19]. We study the sensitivity of symmetry verification to different noise channels, and demonstrate how it can be optimized by adding new symmetries and rotating existing symmetries to be more sensitive to local noise.

II. DEFINITIONS

In this section we cover some basic definitions to be used throughout this paper, and some details of quantum computing that may be skipped by the experienced reader. We use the Pauli group on N qubits $\mathbb{P}^N = \{\mathbb{I}, X, Y, Z\}^{\otimes N}$ throughout this work. These operators form a basis for operators $\hat{O} \in \mathbb{C}^{2^N}$,

$$\hat{O} = \sum_{\hat{P} \in \mathbb{P}^N} O_{\hat{P}} \hat{P}, \quad (1)$$

with coefficients $O_{\hat{P}} \in \mathbb{C}$. If we choose $O_{\hat{P}} \in \mathbb{R}$, this is then a basis for the $2^N \times 2^N$ Hermitian matrices. We call such a basis decomposition of an operator $\hat{O}$ a Pauli decomposition. Such a basis is orthogonal in the Frobenius norm:

$$\|\hat{O}\|_F = \sqrt{\text{Tr}[\hat{O}^\dagger \hat{O}]} \quad (2)$$

Elements $\hat{P} \in \mathbb{P}^N$ have two eigenvalues $p = \pm 1$, with corresponding eigenspaces of dimension 2^{N-1} , and projectors $\hat{M}_p = \frac{1}{2}(1 + p\hat{P})$ onto said eigenspaces. The Pauli group has an additional property; if $\hat{P} \neq \mathbb{I}^{\otimes N} \in \mathbb{P}^N$, $\hat{P}$ commutes with half of the elements of $\mathbb{P}^N$ ($[\hat{P}, \hat{Q}] = 0$), and anticommutes with the other elements ($\{\hat{P}, \hat{R}\} = 0$). This property can be extended to a general operator $\hat{O}$ — $[\hat{P}, \hat{O}] = 0$ ($\{\hat{P}, \hat{O}\} = 0$) if and only if $\hat{P}$ commutes (anticommutes) with each element of the Pauli decomposition of $\hat{O}$ [Eq. (1)].

A quantum computation consists of multiple experiments, each of which can be split into preparations, transformations, and measurements. In the ideal case, a preparation generates a quantum state on a register of N qubits, which is

represented as a vector $|\phi\rangle$ in the complex vector space $\mathbb{C}^{2^N}$. Transformation consists of evolving this state to a new state $|\psi\rangle \in \mathbb{C}^{2^N}$, which may be represented via a unitary operator $|\psi\rangle = U|\phi\rangle$ (with $U \in U(\mathbb{C}^{2^N})$, i.e., $UU^\dagger = U^\dagger U = \mathbb{I}$). Measurement consists of observing the quantum state $|\psi\rangle$ along some degree of freedom. The degree of freedom is represented by a projector-valued measurement $\{\hat{M}_i\}$ for each possible observed value i , where $\sum_i \hat{M}_i = \mathbb{I}$, $\hat{M}_i^2 = \hat{M}_i^\dagger = \hat{M}_i$. The observation records one such value i at random with probability p_i following the Born rule,

$$p_i = |\langle\psi|\hat{M}_i|\psi\rangle|^2, \quad (3)$$

and the state of the system collapses into $\hat{M}_i|\psi\rangle/p_i$.

In the presence of noise, the state of a qubit is instead given by a density matrix $\rho \in \mathbb{D}^N$, where $\mathbb{D}^N$ is the set of $2^N \times 2^N$ positive, trace 1 matrices. These are a generalized form of pure quantum states $|\psi\rangle$, which allow for statistical ensembles of pure states (the well-known adage being that preparing $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ is strictly not the same as preparing $|0\rangle$ or $|1\rangle$ with 50% probability). For every pure state $|\psi\rangle$, the corresponding density matrix is the outer product $|\psi\rangle\langle\psi|$, and the expectation value of an operator $\hat{O}$ may be calculated as

$$\langle\psi|\hat{O}|\psi\rangle = \text{Tr}[\hat{O}|\psi\rangle\langle\psi|]. \quad (4)$$

We will use the latter notation throughout this paper, to be consistent with calculating expectation values on mixed states ρ (where the standard bra-ket expectation value is no longer possible). We will distinguish between operators $\hat{O}$ and density matrices ρ by the use of hats. Note that the trace of products of density matrices is also well defined, and has an obvious interpretation as the overlap between the density matrices, as for pure states,

$$\text{Tr}[|\psi\rangle\langle\psi||\phi\rangle\langle\phi|] = \langle\psi|\phi\rangle\langle\phi|\psi\rangle = |\langle\phi|\psi\rangle|^2. \quad (5)$$

Transformations and measurements of density matrices behave differently to those of pure states [23], but we will not need details of this in this work.

A quantum algorithm incurs a cost based on the number of qubits and coherence time required for quantum hardware to execute it. This cost is usually increased by error mitigation protocols that require additional gates or ancilla qubits. However, these are in general low cost compared to the overhead required for full quantum error correction. Indeed, some error mitigation protocols require no additional quantum hardware or circuitry, hence we define them as “zero cost.” Such protocols may require repetition of the algorithm in order to estimate expectation values $\text{Tr}[\hat{O}\rho]$, but this may be offset by parallelizing across multiple quantum devices. This cost metric is then similar to the quantum volume [7] often used to characterize quantum hardware.

III. SYMMETRY VERIFICATION

Our study is motivated by the presence of symmetries in quantum mechanical systems. In such systems, one has a Hamiltonian $\hat{H}$, and is usually interested in studying the properties of ground or low-lying eigenstates of the system. A (unitary) symmetry of a system is a unitary operator $\hat{S}$ that commutes with the Hamiltonian— $[\hat{H}, \hat{S}] = 0$. When this is

true, $\hat{H}$ may be block diagonalized within the eigenspaces of $\hat{S}$. Then, if one were to study eigenstates of $\hat{H}$ on a quantum computer, one may perform such a study entirely within a single target eigenspace $\mathcal{S}$ of $\hat{S}$. In real-world quantum computers, noise may shift the state of the computer outside of the target eigenspace $\mathcal{S}$. By verifying during or at the end of a calculation that the system remains in $\mathcal{S}$, and throwing away results where this is not the case, it is thus possible to make our quantum computation less sensitive to these types of noise.

Verification of a symmetry is performed by measurement and post-selection which is typically performed in the computational basis (the eigenstates $|0\rangle$ and $|1\rangle$ of a single qubit). The Pauli operators $\mathbb{P}^N$ may be rotated into this basis relatively easily (see Sec. IV), and as such are a good class from which to draw symmetry operators. If $\hat{S} \notin \mathbb{P}^N$, but the target eigenspace $\mathcal{S}$ lies within the eigenspace of a Pauli operator $\hat{P}$, then measuring $\hat{P}$ presents a low-cost alternative to measuring $\hat{S}$, though this may provide less error mitigation in the case where the eigenspace of $\hat{P}$ is strictly larger than $\mathcal{S}$. In general, symmetry verification will work with any construction of a projector valued measurement $\{\hat{M}_i\}$ where one projector $\hat{M}_\mathcal{S}$ projects onto the target eigenspace $\mathcal{S}$. We note that phase estimation [24] provides a generic construction for such a measurement, although this is a rather high cost circuit (in particular requiring the ability to apply the symmetry $\hat{U}$ on the quantum computer). This requirement for measurement implies that symmetry verification cannot be extended to antiunitary symmetries (nor to symmetries that anticommute with the Hamiltonian), as these do not lead to eigenspaces that can be projected into.

The projector valued measurement $\{\hat{M}_i\}$ is the more general object for symmetry verification than the symmetry $\hat{S}$. In an arbitrary quantum circuit at an arbitrary time, if we know by any means that the state $|\psi\rangle$ in the absence of error satisfies $\hat{M}_\mathcal{S}|\psi\rangle = |\psi\rangle$, measuring $\{\hat{M}_i\}$ on the noisy state ρ and post-selecting will project to the state,

$$\rho_s = \frac{\hat{M}_\mathcal{S}\rho\hat{M}_\mathcal{S}}{\text{Tr}[\hat{M}_\mathcal{S}\rho]}. \quad (6)$$

Then, we have

$$\text{Tr}[\rho_s|\psi\rangle\langle\psi|] = \frac{\text{Tr}[\rho|\psi\rangle\langle\psi|]}{\text{Tr}[\hat{M}_\mathcal{S}\rho]} \geq \text{Tr}[\rho|\psi\rangle\langle\psi|], \quad (7)$$

and our new state ρ_s has strictly greater overlap with the target $|\psi\rangle$ than the pre-selection ρ (unless $\hat{M}_\mathcal{S}\rho\hat{M}_\mathcal{S} = \rho$, in which case $\rho_s = \rho$). Such a procedure can be immediately extended to multiple operators $\hat{S}_1, \hat{S}_2, \dots$, as long as $[\hat{S}_i, \hat{S}_j] = 0$. (If this is not the case, sequential symmetry verification projects between different eigenspaces, which is inefficient and greatly increases the number of experiments that must be thrown away.) Symmetry verification may also be repeated at multiple points during a quantum circuit, by inserting measurement of $\hat{S}$ in between gates, as long as we expect the state $|\psi(t)\rangle$ to be an eigenstate of $\hat{S}$ at time t during the circuit. We call such protocols “bulk” symmetry verification, as opposed to “final” symmetry verification at the end of an experiment.

FIG. 1. Quantum circuit for ancilla symmetry verification of a symmetry $\hat{S}$. (a) A simple circuit entangling all qubits with a single ancilla qubit. The rotations $\hat{R}_i$ depend on the tensor components $\hat{S}_i$ on each qubit i (relationship given in text). (b) A circuit making an identical measurement to that in (a), but with only local CNOT and SWAP two-qubit gates. A SWAP between qubit 0 and the ancilla is not required because the Bell state prepared after the first CNOT is symmetric between the two qubits (this is not the case for the remaining qubits).

IV. ANCILLA AND IN-LINE SYMMETRY VERIFICATION

The simplest form of the symmetry verification involves the use of an ancilla qubit to measure the Pauli symmetry $\hat{S}$. Let us write $\hat{S} \in \mathbb{P}^n$ in terms of its tensor factors; $\hat{S} = \otimes_i \hat{S}_i$, and let N_S be the number of nontrivial $\hat{S}_i = \{X, Y, Z\}$. To each such $\hat{S}_i$, we can associate a corresponding rotation $\hat{R}_i = \{\exp(i\frac{\pi}{2}Y), \exp(-i\frac{\pi}{2}X), \mathbf{1}\}$ (such that $\hat{R}_i|\hat{S}_i = 1\rangle = |0\rangle$). The verification circuit is then shown in Fig. 1(a). For each nontrivial $\hat{S}_i$, the corresponding qubit is rotated by $\hat{R}_i$, then performs a controlled-NOT gate on the ancilla qubit, and finally is rotated by $\hat{R}_i^{-1}$. This requires that the ancilla qubit be coupled to each qubit in the system register that it measures, which is in general not possible in a quantum circuit. As a low-cost alternative [Fig. 1(b)], the ancilla qubit may be shuffled along the system register via SWAP gates as it performs the controlled phase gate. In either case, as the ancilla qubit must interact with each register qubit individually, the circuit depth must be $O(N_S)$.

It is possible to forego the ancilla qubit in symmetry verification, by instead encrypting the symmetry $\hat{S}$ onto the computational degree of freedom of a qubit within the system itself, which is then read out. In Fig. 2(a) we give an example circuit for this in-line symmetry verification, with circuit depth only $O(\log(N_S))$. This logarithmic depth requires qubits to be coupled as a binary tree, which is not possible in systems which allow only local couplings. In general, for such a d -dimensional local coupling, the depth of the circuit must be at least $O(N_S^{1/d})$, being the minimum depth of a light-cone encompassing N_S qubits. In Fig. 2(b) we give such a circuit for a system with linear connectivity. Even when all-to-all coupling is available, the $O(\log(N_S))$ -depth circuit [Fig. 2(a)]

FIG. 2. Quantum circuits for in-line symmetry verification. (a) The optimal verification circuit has $O(\log(N_S))$ depth, but requires long-range connectivity between qubits, which is not available on many architectures. (b) In the presence of linear connectivity, an $O(N_S)$ depth verification circuit is optimal.

may not be preferable, as the duty cycle for each qubit (i.e., the period of time between the first and last gate each qubit is involved in) is length $O(\log(N_S))$. By contrast, the duty cycle of an individual qubit during the circuit in Fig. 1(b) is $O(1)$. A short duty cycle implies that qubits can be used to perform other operations while the symmetry verification is ongoing, reducing the time cost when this circuit is performed as a small block of a larger computation.

V. VARIATIONAL QUANTUM EIGENSOLVERS

As an example target algorithm for symmetry verification, we consider ground-state preparation for a Hamiltonian $\hat{H}$ via a variational quantum eigensolver [9,25]. An (ideal) variational quantum eigensolver consists of a unitary circuit $U(\vec{\theta})$, parametrized by a vector of free angles $\vec{\theta}$ that control individual gates within the circuit. This circuit acts on a starting state, which we take to be the computational basis state $|0, \dots, 0\rangle$, to produce a variational final state $|\psi(\vec{\theta})\rangle = U(\vec{\theta})|0, \dots, 0\rangle$. These angles are controlled classically to minimize the energy $E(\vec{\theta}) = \langle\psi(\vec{\theta})|\hat{H}|\psi(\vec{\theta})\rangle$. This expectation value is calculated in an experiment by taking the Pauli decomposition of $\hat{H}$ [Eq. (1)], preparing $|\psi(\vec{\theta})\rangle$ and measuring each $\hat{P}_i$ repeatedly to accumulate statistics on $\langle\psi|\hat{P}_i|\psi\rangle$.

Variational quantum eigensolvers (VQE) are natural candidates for final symmetry verification, and common classes of VQEs are also natural candidates for bulk symmetry verification. In particular, for fermionic systems (such as the electronic structure problem), global fermion parity is conserved, making it a prime target for symmetry verification. (At low energy, for nonsuperconducting systems, the particle number is often conserved as well, but this is not a Pauli operator, and is much more difficult to measure.) Using the Jordan-

Wigner transformation on an N -fermion Hamiltonian, this symmetry takes the form $Z^{\otimes N}$. Most VQEs consist of creating an approximate starting state (such as the Hartree-Fock state) that respects this symmetry, and then performing multiple local rotations that continue to respect this symmetry. This is true of both the unitary coupled cluster (UCC) ansatz [9], and the quantum approximate optimization algorithm (QAOA) [26]. In the former, the ansatz is taken as the expansion of the cluster operator $e^{\hat{T}-\hat{T}^\dagger}$,

$$\hat{T} = \sum_n \hat{T}^{(n)}, \quad (8)$$

$$\hat{T}^{(n)} = \sum_{i_1, \dots, i_n; j_1, \dots, j_n} \theta_{j_1, \dots, j_n}^{i_1, \dots, i_n} \left(\prod_{m=1}^n \hat{c}_{i_m}^\dagger \right) \left(\prod_{m=1}^n \hat{c}_{j_m} \right), \quad (9)$$

where the θ parameters are taken as the free parameters to be optimized, and the sum is a sum over empty molecular orbitals to the left of the semicolon, and filled molecular orbitals to the right. This exponentiation is typically performed by the Trotter-Suzuki expansion, leaving a series of unitaries,

$$\prod_{i,j} e^{\theta_{ij}^{ij} (\hat{c}_i^\dagger \hat{c}_j - \hat{c}_j^\dagger \hat{c}_i)} \prod_{i,j;k,l} e^{\theta_{ijkl}^{ijkl} (\hat{c}_i^\dagger \hat{c}_j \hat{c}_k \hat{c}_l - \hat{c}_l^\dagger \hat{c}_k \hat{c}_j^\dagger \hat{c}_i)} \dots, \quad (10)$$

each of which respects fermion parity. QAOA for the electronic structure problem consists of performing steps of time evolution alternating between the Hartree-Fock Hamiltonian and the electronic-structure Hamiltonian, both of which respect fermion parity. Thus, for both ansatz, bulk symmetry verification could be performed between individual steps of the time evolution.

Although symmetry verification promises a final state with greater overlap with the ground state, it does not promise a necessarily lower energy. Let us write the (un-normalized) symmetry-accepted state ρ_s , and the symmetry-rejected state ρ_r . If our measurement was perfect, we would have

$$\rho_s = \hat{M}_s \rho \hat{M}_s, \quad \rho_r = (\mathbb{I} - \hat{M}_s) \rho (\mathbb{I} - \hat{M}_s). \quad (11)$$

Then, $\text{Tr}[\hat{H} \rho] = \text{Tr}[\hat{H} \rho_r] + \text{Tr}[\hat{H} \rho_s]$. Now, suppose the rejected state ρ_r has lower energy than the accepted state ρ_s ;

$$\frac{\text{Tr}[\hat{H} \rho_r]}{\text{Tr}[\rho_r]} < \frac{\text{Tr}[\hat{H} \rho_s]}{\text{Tr}[\rho_s]}. \quad (12)$$

We can calculate

$$\begin{aligned} \text{Tr}[\hat{H} \rho] &= \text{Tr}[\rho_r] \frac{\text{Tr}[\hat{H} \rho_r]}{\text{Tr}[\rho_r]} + \text{Tr}[\rho_s] \frac{\text{Tr}[\hat{H} \rho_s]}{\text{Tr}[\rho_s]} \\ &< (\text{Tr}[\rho_r] + \text{Tr}[\rho_s]) \frac{\text{Tr}[\hat{H} \rho_s]}{\text{Tr}[\rho_s]} = \frac{\text{Tr}[\hat{H} \rho_s]}{\text{Tr}[\rho_s]}, \end{aligned}$$

and our symmetry-verified state would be higher in energy than the initial state as well. As the energy of ρ_r lies strictly above the ground state, failure of symmetry verification must imply ρ_s has sufficiently large overlap with high-energy states. As such, we would suggest that such a failure implies the energy of ρ itself is not to be trusted.

VI. POST-SELECTED SYMMETRY VERIFICATION AND S-QSE

Conveniently, when a quantum computation requires calculating the expectation values of a set of Pauli operators, symmetry verification may be performed via post-processing of the expectation values themselves (with possibly some additional measurements), rather than requiring additional quantum circuitry. Suppose we want to calculate the expectation value of $\hat{P} \in \mathbb{P}^N$ on our state ρ following projection onto the $\hat{S} = s (= \pm 1)$ subspace of our symmetry $\hat{S} \in \mathbb{P}^N$. The projector onto this subspace may be written $\hat{M}_s = \frac{1}{2}(\mathbb{I} + s\hat{S})$. Then, the expectation value of $\hat{P}$ on the state ρ_s targeted by the symmetry verification can be expanded using Eq. (6),

$$\begin{aligned} \text{Tr}[\hat{P} \rho_s] &= \text{Tr} \left[\hat{P} \frac{\hat{M}_s \rho \hat{M}_s}{\text{Tr}[\hat{M}_s \rho]} \right] \\ &= \frac{\text{Tr}[\hat{P} \rho] + s \text{Tr}[\hat{P} \hat{S} \rho]}{1 + s \text{Tr}[\hat{S} \rho]}, \end{aligned} \quad (13)$$

where we have used the cyclic property of the trace and the fact that $[\hat{P}, \hat{M}_s] = 0$ to write $\text{Tr}[\hat{P} \hat{M}_s \rho \hat{M}_s] = \text{Tr}[\hat{P} \hat{M}_s \rho]$, and expanded our definition of $\hat{M}_s$. The expectation values $\text{Tr}[\hat{S} \rho]$, $\text{Tr}[\hat{P} \rho]$, and $\text{Tr}[\hat{P} \hat{S} \rho]$ may be then calculated using the unverified state ρ , and substituted into Eq. (13) to obtain the verified result. By avoiding additional quantum circuitry, we expect this method to outperform both ancilla and in-line symmetry verification. However, we note that post-selection cannot be used for bulk symmetry verification (as we cannot measure these expectation values during the circuit). Furthermore, it cannot be used in algorithms where the output is not an expectation value $\text{Tr}[\hat{P} \rho]$.

Post-selected symmetry verification can be observed to be identical to a form of the quantum subspace expansion (QSE) [19]. Originally designed to investigate the low-energy excited states around the ground space found by a variational quantum eigensolver, QSE works by taking a set of excitation operators $\{\hat{E}_i\}$, which can be applied to the approximated ground state $|\psi(\theta)\rangle$ to obtain a set of states $|\phi_i\rangle = \hat{E}_i |\psi(\theta)\rangle$. The spectrum of the Hamiltonian within the manifold spanned by these states can be calculated as the solution to the generalized eigenvalue problem,

$$\hat{H}_{\text{QSE}} |\xi\rangle = \lambda \hat{B}_{\text{QSE}} |\xi\rangle. \quad (14)$$

Here, $\hat{H}_{\text{QSE}}$ is the Hamiltonian matrix projected into the spanned manifold,

$$[\hat{H}_{\text{QSE}}]_{i,j} = \text{Tr}[\hat{H} |\phi_i\rangle \langle \phi_j|], \quad (15)$$

and $\hat{B}_{\text{QSE}}$ is the overlap matrix,

$$[\hat{B}_{\text{QSE}}]_{i,j} = \text{Tr}[|\phi_i\rangle \langle \phi_j|], \quad (16)$$

to account for the fact that $|\phi_i\rangle$ and $|\phi_j\rangle$ are in general not orthogonal. In the presence of noise, although the state $|\phi_i\rangle$ is not well defined (as our noisy state ρ is not a pure state), the operators $|\phi_i\rangle \langle \phi_j| = \hat{E}_i \rho \hat{E}_j^\dagger$ remain well defined, and the expectation values in Eqs. (15) and (16) are still able to be measured in an experiment.

The set $\{\hat{E}_i\}$ is usually taken to be the set of low-order polynomials in qubit or fermion operators [19,20]. However,

if the set $\{\mathbb{I}, \hat{S}\}$ is chosen as excitation operators, the solution to the generalized eigenvalue problem is the same as that obtained by post-selection. To show this, we expand

$$\begin{aligned}\text{Tr}[\hat{H}\rho_s] &= \sum_i h_i \text{Tr}[\hat{P}_i \rho_s] \\ &= \sum_i \frac{\text{Tr}[h_i \hat{P}_i \rho] + s \text{Tr}[h_i \hat{P}_i \hat{S} \rho]}{1 + s \text{Tr}[\hat{S} \rho]} \\ &= \frac{\text{Tr}[\hat{H} \rho] + s \text{Tr}[\hat{H} \hat{S} \rho]}{1 + s \text{Tr}[\hat{S} \rho]}.\end{aligned}\quad (17)$$

Next, we calculate the QSE matrices (using the commutation of $\hat{H}$ and $\hat{S}$),

$$\hat{H}_{\text{QSE}} = \begin{bmatrix} \text{Tr}[\hat{H} \rho] & \text{Tr}[\hat{H} \hat{S} \rho] \\ \text{Tr}[\hat{H} \hat{S} \rho] & \text{Tr}[\hat{H} \rho] \end{bmatrix}, \quad (18)$$

$$\hat{B}_{\text{QSE}} = \begin{bmatrix} 1 & \text{Tr}[\hat{S} \rho] \\ \text{Tr}[\hat{S} \rho] & 1 \end{bmatrix}. \quad (19)$$

Assuming that $\text{Tr}[\hat{S} \rho] \neq 1$, $\hat{B}_{\text{QSE}}$ is invertible, the problem reduces to finding the (regular) eigenvalues of

$$\hat{B}_{\text{QSE}}^{-1} \hat{H}_{\text{QSE}} = \frac{1}{1 - \text{Tr}[\hat{S} \rho]^2} \begin{bmatrix} \alpha & \beta \\ \beta & \alpha \end{bmatrix}, \quad (20)$$

where

$$\alpha = \text{Tr}[\hat{H} \rho] - \text{Tr}[\hat{S} \rho] \text{Tr}[\hat{H} \hat{S} \rho], \quad (21)$$

$$\beta = \text{Tr}[\hat{H} \hat{S} \rho] - \text{Tr}[\hat{H} \rho] \text{Tr}[\hat{S} \rho]. \quad (22)$$

The eigenvalues of this matrix take the form,

$$\lambda = \frac{1}{1 - \text{Tr}[\hat{S} \rho]^2} (\alpha \pm \beta) \quad (23)$$

$$= \frac{\text{Tr}[\hat{H} \rho] \pm \text{Tr}[\hat{H} \hat{S} \rho]}{1 \pm \text{Tr}[\hat{S} \rho]}, \quad (24)$$

which can be seen to be equal to those found in Eq. (17). We call this version of the quantum subspace expansion symmetry-QSE, or S-QSE for short.

This result is not surprising; it was suggested in [19] to account for symmetries during QSE by projecting $\hat{H}_{\text{QSE}}$ and $\hat{B}_{\text{QSE}}$ into the symmetry subspace, which achieves the same result as in the above. However, this demonstrates that one may account for symmetries via a version of QSE without calculating the full linear response. Moreover, this implies that S-QSE corrects for both coherent and incoherent errors that project out of the $\hat{S} = s$ subspace. By contrast, QSE with an operator that anticommutes with the Hamiltonian can only correct coherent errors (see Appendix). S-QSE may be immediately combined with other forms of QSE, for example, linear response QSE, by including both sets of operators as excitations.

VII. SIMULATION OF SYMMETRY VERIFICATION ON THE HYDROGEN MOLECULE

To first investigate symmetry verification in a simple setting, we use a VQE to find the ground-state energy of H_2

on two qubits. This follows previous experimental demonstrations [9,20,27,28]. We take the STO-3G basis for H_2 , which has four spin orbitals, and convert this into a qubit Hamiltonian via the Bravyi-Kitaev transformation. The four spin orbitals require four qubits to represent them on, but in this representation the Hamiltonian is diagonal on two of the qubits, which may be removed. The remaining two-qubit Hamiltonian takes the form,

$$\hat{H} = h_0 \mathbb{I} + h_1 \mathbb{I} Z + h_2 Z \mathbb{I} + h_3 X X + h_4 Y Y + h_5 Z Z, \quad (25)$$

where h_i are sums of integrated two- and four-body terms from the original electronic structure problem. The calculation of these terms, and the Bravyi-Kitaev transformation itself, were performed using the PSI4 [29] and OPENFERMION [30] packages. The Hamiltonian can be seen to commute with the symmetry $\hat{S} = ZZ$. Our ground-state wave function has nontrivial overlap with the Hartree-Fock wave function, which is in the $ZZ = -1$ subspace; this is then our target subspace. We follow the unitary coupled cluster ansatz of [27], which consists of exciting our system to the $|01\rangle$ state, and performing the unitary rotation,

$$\hat{U}(\theta) = e^{-i\theta X_0 Y_1}. \quad (26)$$

This unitary rotation may be decomposed using standard methods [31]. As described previously, the VQE procedure consists of fixing θ , repeatedly preparing $|\psi(\theta)\rangle$ and measuring collections of terms in the Pauli decomposition of $\hat{H}$ until a good estimate of the energy $E(\theta)$ is found. This is then repeated at varying θ as demanded by a classical optimizer until a minimum $E(\theta)$ is found [9].

We compare the performance of the three symmetry verification protocols described previously as a final symmetry verification step. The ancilla symmetry verification is performed in the same manner as Fig. 1(a). The in-line symmetry verification is performed in a manner similar to Fig. 2(a), but as this is final symmetry verification, we have no need to undo the symmetry measurement. Instead, to measure the expectation value of a Pauli operator $\text{Tr}[\rho \hat{P}]$, we can propagate $\hat{P}$ through the symmetry verification circuit [32] and measure the corresponding Pauli term. It is then sufficient to rotate the control qubit to recover the expectation values $\langle ZZ \rangle$ and $\langle XX \rangle$. From this we may calculate all other expectation values in Eq. (25) using the fact that $ZZ = -1$. For this problem, S-QSE not only requires no additional circuitry, but also no additional measurements (all required terms are in the Pauli decomposition of the Hamiltonian).

To test symmetry verification in the presence of realistic noise, we simulate our chosen experiment using the quantum-sim density matrix simulator [33]. We take gate error models and parameters similar to previous simulation work based on experimental data of state-of-the-art superconducting transmon qubits [34]. Errors in transmon qubits are dominated by decoherence times, which we take at a base level to be $T_1 = T_2 = 20 \mu\text{s}$. This should be compared to single- and two-qubit gate times of 20 ns (giving a total circuit length without symmetry verification of 220 ns). Single- and two-qubit gates suffer from additional dephasing noise of 0.01% and 1%, respectively. We assume that single-shot measurement (for

FIG. 3. Accuracy of the VQE over the entire bond dissociation curve using the different symmetry verification methods mentioned in the text (labeled in legend). (Top) The target curve of H_2 , compared to the exact result (black line). (Bottom) Log plot of the difference between the black lines and points in the above plot.

verification purposes) has a read-out error of 1%, and that error in tomographic measurements and pre-rotations (used to calculate the expectation values themselves) can be canceled by linear inversion tomography [35,36].

Using the above error model, we observe (Fig. 3) that the unmitigated VQE (blue points) achieves an error in the energy of approximately 0.01–0.04 hartree across the bond dissociation curve. This error is improved upon by all symmetry verification techniques. S-QSE (red diamonds) provides the largest improvement of all symmetry verification protocols, as no additional errors are introduced. The S-QSE circuit is observed to give approximately a fivefold improvement over the unmitigated circuit, while ancilla (orange crosses) and in-line (green squares) symmetry verification show an approximately twofold and threefold improvement, respectively. The differences between S-QSE and other forms of symmetry verification emphasize the importance of minimizing the verification cost in bulk symmetry verification (where S-QSE is no longer available).

We now investigate the effect of different noise channels on the performance of symmetry verification. Any noise channel that commutes with the symmetry operators evolves the system state within the target subspace, which symmetry verification explicitly does not mitigate. The analysis of which channels have this property can be reduced to an analysis over $\mathbb{P}^N$, as if we mitigate Pauli errors $\hat{P}_i \in \mathbb{P}^N$, we also mitigate any linear combination of them [21]. In the above circuit, the ZZ symmetry commutes with any single-qubit Z errors, making the protocol prone to the T_ϕ (pure dephasing) channel, but it anticommutes with single-qubit X errors, making the protocol resilient against the T_1 (amplitude decay) channel. To investigate this, in Fig. 4 we calculate the error in determining the ground-state energy near the minima of the bond dissociation curve (0.75 Å bond distance) using S-QSE, as we vary T_1 and T_ϕ . We turn all other error sources off, and vary T_1 (T_ϕ) with $T_\phi = 20 \mu s$ ($T_1 = 20 \mu s$) fixed. In the absence of error

FIG. 4. Effect of varying decoherence times on the VQE accuracy. With all other error sources turned off, T_1 is varied with $T_\phi = 20 \mu s$ fixed (red-dashed curves), and T_ϕ is varied with $T_1 = 20 \mu s$ fixed (blue-solid curves). We plot the error in estimating the ground-state energy for the unmitigated experiment (squares), and the circuit mitigated with S-QSE (circles). Data points for the blue and red curves are identical at $T_1 = T_\phi = 20 \mu s$, as can be seen from the complete overlap.

mitigation, the two decoherence sources have almost identical effect (deviation approximately 10^{-2} hartree). However, in the presence of error mitigation, the susceptibility of the VQE to T_1 noise is noticeably smaller than to T_ϕ noise—up to a factor of two over the range of decoherence times plotted. We note that S-QSE does not make our circuit second-order sensitive to T_1 noise. This can be understood as X errors at some points during our VQE circuit are rotated to Z errors by later gates in the circuit, preventing their mitigation.

VIII. INSERTING AND ROTATING SYMMETRIES

As observed in the previous section, verifying single symmetries has a marked effect on the performance of a quantum circuit, but will not catch and remove all sources of noise. In this section we suggest how one may improve upon this by adding additional symmetries to the quantum algorithm, and by rotating existing symmetries to make them more sensitive to errors on the underlying quantum hardware. In the language of quantum error correction, this is a low-cost attempt to increase the distance of the detection code.

We first suggest a method to extend an N -qubit Hamiltonian $\hat{H}$, given a Pauli operator $\hat{P} \in \mathbb{P}^N$, to an $N + 1$ -qubit Hamiltonian $\hat{H}_{\text{ext}}$,

$$\hat{H}_{\text{ext}} = \begin{bmatrix} \hat{H} & 0 \\ 0 & \hat{P}\hat{H}\hat{P} \end{bmatrix}. \quad (27)$$

Both blocks of $\hat{H}_{\text{ext}}$ can be seen to have the same eigenspectrum (as this is unaffected by the unitary rotation of $\hat{P}$), and $\hat{H}_{\text{ext}}$ commutes with the operator,

$$\begin{bmatrix} 0 & \hat{P} \\ \hat{P} & 0 \end{bmatrix} = X\hat{P}, \quad (28)$$

which is then the new symmetry operator. This mapping corresponds to mapping Pauli operators $\hat{Q} \in \mathbb{P}^N$ in the original problem to

$$\hat{Q}_{\text{ext}} = \begin{cases} \mathbb{I}\hat{Q} & \text{if } [\hat{Q}, \hat{P}] = 0 \\ X\hat{Q} & \text{if } \{\hat{Q}, \hat{P}\} = 0 \end{cases}. \quad (29)$$

To implement this in the algorithm itself, we note that every circuit can be decomposed into a product of unitary rotations,

$$\prod_j e^{i\theta_j \hat{Q}_j} \quad \hat{Q}_j \in \mathbb{P}^N, \quad (30)$$

where a single $\hat{Q} \in \mathbb{P}^N$ may be repeated in the product. Adding the symmetry then consists of replacing these rotations by rotations around the transformed operator $\hat{Q}_{\text{ext}}$ [as per Eq. (29)], and re-decomposing the operations into a circuit (using, e.g., the methods of [31,37]). If $\hat{H}$ had a previous set of symmetries $\hat{S}_i$, these are transformed to a new set $\hat{S}_{i,\text{ext}}$ [following Eq. (29)], that commute with both $\hat{H}_{\text{ext}}$ and the additional symmetry $X\hat{P}$. This extension method is particularly suitable for digital quantum simulation, as circuits are often generated in the form of Eq. (30). This is the case for traditional Hamiltonian simulation [38], quantum phase estimation [31], and the UCC QSE discussed previously, all of which require exponentiating an operator via the Suzuki-Trotter expansion [39].

Beyond choosing the number of symmetries in a problem, one may wish to choose how these symmetries appear in the problem. In particular, sets of symmetries may be found that anticommute with all local operators, which should increase the mitigation power of the verification protocol against local sources of noise. (For example, the N -qubit operators $X^{\otimes N}$ and $Z^{\otimes N}$ with even N .) Any two groups of M Pauli operators are unitarily equivalent as long as they satisfy the same commutation and multiplication rules (e.g., $\mathbb{I}Z$, $Z\mathbb{I}$, and ZZ are equivalent to XX , YY , and ZZ , but not to $\mathbb{I}X$, $\mathbb{I}Y$, and $\mathbb{I}Z$). To find such unitary transformations, we suggest decomposing them into rotations of the form $\hat{R} = e^{i\frac{\pi}{2}\hat{Q}}$ for $\hat{Q} \in \mathbb{P}$, which transforms

$$\hat{P} \in \mathbb{P} \rightarrow \hat{R}^\dagger \hat{P} \hat{R} = \begin{cases} \hat{P} & \text{if } [\hat{P}, \hat{Q}] = 0 \\ i\hat{P}\hat{Q} & \text{if } \{\hat{P}, \hat{Q}\} = 0 \end{cases} \quad (31)$$

Rotations of this form have a few desirable properties. Their effect is easy to calculate classically, and they transform Pauli operators to Pauli operators. Furthermore, each $\hat{R}$ leaves half of the Pauli group unchanged. This allows for some choice of rotations to leave desired symmetries (or other operators) already present in the problem invariant, while other terms are rotated.

IX. EXTENDING THE SYMMETRY VERIFICATION OF THE HYDROGEN MOLECULE

We now demonstrate the verification of multiple symmetries by extending the previous VQE simulation of H_2 . We transform the electronic structure Hamiltonian onto a qubit representation this time via the Jordan-Wigner transformation. This gives the four-qubit Hamiltonian,

$$\hat{H} = h_I \mathbb{I} + \sum_i h_i Z_i + \frac{1}{2} \sum_{i \neq j} h_{i,j} Z_i Z_j + h_s (X_0 Y_1 Y_2 X_3 + Y_0 X_1 X_2 Y_3 - X_0 X_1 Y_2 Y_3 - Y_0 Y_1 X_2 X_3), \quad (32)$$

which has symmetries $\hat{S}_0 = Z_0 Z_1$, $\hat{S}_1 = Z_0 Z_2$, and $\hat{S}_2 = Z_0 Z_1 Z_2 Z_3$. In the Bravyi-Kitaev transformation these symmetries were the additional qubits that were thrown away. We

FIG. 5. Adding and adjusting symmetries to optimize symmetry verification. The blue (dots) and red (diamonds) curves correspond to their colored (shaped) counterparts in Fig. 3, while the purple (squares) and brown (crosses) curves come from a four-qubit simulation of H_2 using the two protocols described in the text. The dashed lines represent the S-QSE versions of their solid counterparts. Error parameters on all qubits are the same for all simulations (parameters given in the text).

choose again the unitary coupled cluster ansatz for the VQE, which can be reduced to the operator [40],

$$\hat{U}(\theta) = e^{i\theta Y_0 X_1 X_2 X_3}. \quad (33)$$

As in the two-qubit case, the VQE circuit consists of preparing the system in the Hartree-Fock state $|1100\rangle$, applying $U(\theta)$ and measuring the variational energy, for a total circuit time of 400 ns.

The above set of symmetries still commute with all single-qubit Z errors, so we rotate our problem to increase the mitigation power of symmetry verification. We choose the rotation,

$$\hat{R} = e^{i\frac{\pi}{2} Y_0 X_2} e^{i\frac{\pi}{2} Y_1 X_3}. \quad (34)$$

This transforms the symmetry $\hat{S}_0 \rightarrow X_0 X_1 X_2 X_3$, while leaving $\hat{S}_1$ and $\hat{S}_2$ unchanged. The resulting set of symmetries do not commute with any single-qubit X or Z operator, as required. To create the transformed circuit, we need to transform both our starting state $|1100\rangle \rightarrow \hat{R}|1100\rangle$, and the UCC unitary ansatz,

$$\hat{U}(\theta) \rightarrow \hat{R} \hat{U} \hat{R}^\dagger = e^{i\theta Y_0 Z_1 X_2}. \quad (35)$$

The transformed circuit incurs an additional cost from this initial application of $\hat{R}$, but this is balanced by the reduced weight of the transformed cluster operator, resulting in a total circuit time of 440 ns.

In Fig. 5, we compare the performance of the two different circuits above to the two-qubit circuit of Fig. 3, with and without the addition of S-QSE. At small bond distance ($\lesssim 0.75$ Å), the target ground state (in the absence of rotation by $\hat{R}$) is roughly a computational basis state, which is immune to dephasing errors. At this point, all three verification protocols perform roughly similarly, despite the unmitigated four-qubit simulations performing significantly worse than the unmitigated two-qubit simulation. At large bond distance ($\gtrsim 0.75$ Å), the ground state is prone to T_2 noise, at which point we see the rotated four-qubit S-QSE simulation significantly outperforming its counterparts. At the largest distance

studied, this simulation achieves a twofold reduction in error compared to the two-qubit S-QSE simulation, despite using twice as many qubits and a twice as long circuit. By comparison, unrotated S-QSE on four qubits cannot protect against the T_2 noise accumulated over the simulation, and performs a factor of two worse than the two-qubit S-QSE simulation. This clearly demonstrates the need to optimize symmetry verification protocols to account for errors present in the system as this technique is scaled up to larger computations. Over the entire bond-dissociation curve, the rotated four-qubit S-QSE simulation outperforms its unmitigated counterpart by over an order of magnitude.

X. CONCLUSION

In this paper we have presented a low-cost strategy for error mitigation, which we call symmetry verification. We have discussed various ways in which it can be applied to different algorithms, and various methods to optimize the mitigation power against common sources of error. We have demonstrated these protocols on a simulated VQE experiment of H_2 , and observed that they outperform the unmitigated result over the entire bond-dissociation curve by around an order of magnitude.

Although the above techniques are very promising for small experiments, much work needs to be done optimizing symmetry verification for midrange experiments in the NISQ era. The addition and choice of symmetries needs to be investigated further to minimize the resulting circuit depth. Further study is also needed on the optimal number of symmetry verifications to be added to a circuit, both to maximize mitigation and minimize run time (which increases exponentially in the number of verifications made). Finally, given the obvious connection between symmetry verification and the stabilizer formalism of quantum error correction, it is natural to ask whether one can mix the two to transform slowly between midsize NISQ circuits and large-scale fault-tolerant ones.

While this paper was in production, a related work by McArdle *et al.* [41] appeared on the ArXiv. They simulate the performance of ancilla symmetry verification for a VQE, and its combination with other error mitigation strategies to further improve robustness against noise. Their results are consistent with and complementary to our own, and they provide useful techniques for measuring non-Pauli operators not considered in this work.

ACKNOWLEDGMENTS

The authors wish to thank Carlo Beenakker, Leonardo DiCarlo, Brian Tarasinski, Barbara Terhal, Detlef Hohl, Luuk Visscher, Francesco Buda, Yaroslav Herasymenko, and Adriaan Rol for feedback, advice, and support in this project. This research was supported by the Netherlands Organization for Scientific Research (NWO/OCW) and by an ERC Synergy Grant.

APPENDIX: ERROR MITIGATION OF QSE WITH ANTICOMMUTING OPERATORS

In this appendix we repeat the analysis of QSE from the text, but with an operator $\hat{A}$ that anticommutes with the

Hamiltonian $\hat{H}$. Let us assume to begin that $\hat{A}$ is unitary. Such an operator cannot be simultaneously diagonalized with $\hat{H}$, and so we have no result from symmetry verification to compare with. Given an eigenstate $\hat{H}|\psi\rangle = E|\psi\rangle$, we have that $\hat{H}\hat{A}|\psi\rangle = -\hat{A}\hat{H}|\psi\rangle = -E\hat{A}|\psi\rangle$, and so the presence of an anticommuting operator splits the eigenstates of $\hat{H}$ into pairs of equal magnitude but opposite sign energies (known as eigenstates of different chirality). If $\hat{A} = \hat{A}^\dagger$, the eigenstates of $\hat{A}$ itself are the equal superpositions,

$$|\pm\rangle = \frac{1}{\sqrt{2}}(|\psi\rangle \pm \hat{A}|\psi\rangle). \quad (A1)$$

For QSE, we must calculate the operators $\hat{H}_{\text{QSE}}$ and $\hat{B}_{\text{QSE}}$.

$$\hat{B}_{\text{QSE}} = \begin{bmatrix} 1 & \text{Tr}[\hat{A}\rho] \\ \text{Tr}[\hat{A}^\dagger\rho] & 1 \end{bmatrix}. \quad (A2)$$

$$\hat{H}_{\text{QSE}} = \begin{bmatrix} \text{Tr}[\hat{H}\rho] & \text{Tr}[\hat{H}\hat{A}\rho] \\ \text{Tr}[-\hat{H}\hat{A}^\dagger\rho] & -\text{Tr}[\hat{H}\rho] \end{bmatrix}. \quad (A3)$$

Again assuming $|\text{Tr}[\hat{A}\rho]|^2 \neq 1$, we can invert $\hat{B}_{\text{QSE}}$ and calculate

$$\hat{B}_{\text{QSE}}^{-1} \hat{H}_{\text{QSE}} = \frac{1}{1 - |\text{Tr}[\hat{A}\rho]|^2} \begin{bmatrix} \alpha & \beta \\ -\beta^* & -\alpha^* \end{bmatrix}, \quad (A4)$$

where

$$\alpha = \text{Tr}[\hat{H}\rho] + \text{Tr}[\hat{A}\rho]\text{Tr}[\hat{H}\hat{A}\rho], \quad (A5)$$

$$\beta = \text{Tr}[\hat{H}\hat{A}\rho] + \text{Tr}[\hat{H}\rho]\text{Tr}[\hat{A}\rho]. \quad (A6)$$

The solution to the equation is

$$E_{\text{QSE}}^2 = \frac{|\alpha|^2 + |\beta|^2}{(1 - |\text{Tr}[\hat{A}\rho]|^2)^2} \quad (A7)$$

$$= \frac{\text{Tr}[\hat{H}\rho]^2 + |\text{Tr}[\hat{H}\hat{A}\rho]|^2}{1 - |\text{Tr}[\hat{A}\rho]|^2}. \quad (A8)$$

To understand the gain in energy, $|\text{Tr}[\hat{H}\hat{A}\rho]|^2$, let us first consider a single set of opposite chirality states $|\psi\rangle$ and $\hat{A}|\psi\rangle$ (with energy $\pm E$). We first note that if ρ is an incoherent superposition of the eigenstates,

$$\rho = |a|^2|\psi\rangle\langle\psi| + |b|^2\hat{A}|\psi\rangle\langle\psi|\hat{A}, \quad (A9)$$

$\text{Tr}[\hat{H}\hat{A}\rho] = \text{Tr}[\hat{A}\rho] = 0$ (as $\langle\psi|\hat{A}|\psi\rangle = 0$), and QSE strictly does not improve on the estimate of the ground-state energy. We next consider the opposite situation, where ρ is a coherent superposition of eigenstates:

$$\rho = (\cos(\theta)|\psi\rangle + \sin(\theta)e^{i\phi}\hat{A}|\psi\rangle) \times (\cos(\theta)\langle\psi| + \sin(\theta)e^{-i\phi}\langle\psi|\hat{A}^\dagger). \quad (A10)$$

We can calculate

$$\text{Tr}[\hat{H}\rho] = E \cos(2\theta), \quad (A11)$$

$$\text{Tr}[\hat{A}\rho] = \sin(2\theta)(1 + Ae^{i\phi}), \quad (A12)$$

$$\text{Tr}[\hat{H}\hat{A}\rho] = E \sin(2\theta)(Ae^{i\phi} - 1), \quad (A13)$$

where $A = \langle \psi | \hat{A}^2 | \psi \rangle$ (so $|A| \leq 1$, and for $\hat{A} \in \mathbb{P}^N$, $A = 1$). This gives

$$E_{\text{QSE}}^2 = E^2 \frac{\cos^2(2\theta) + \sin^2(2\theta)\chi_+}{1 - \sin^2(2\theta)\chi_-}, \quad (\text{A14})$$

$$\chi_{\pm} = (1 \pm Ae^{i\phi})(1 \pm Ae^{-i\phi}). \quad (\text{A15})$$

We see that if $A = 1$, $\phi = \frac{\pi}{2}$, QSE corrects the coherent error entirely, while if $A = 1$, $\phi = 0$ it has no effect. This implies that QSE cannot correct coherent rotations of ρ from $|\psi\rangle$ towards an eigenstate of $\hat{A}$. This is in keeping with the general observations in [19] for the performance of QSE as an error mitigation strategy.

If $\hat{A}$ is not unitary, then $\hat{A}^\dagger \hat{A}$ is a Hermitian operator that commutes with $\hat{H}$. Importantly, if $\{\hat{A}, \hat{H}\} = 0$, $\{\hat{A}^\dagger \hat{A}, \hat{H}\} = 0$ as well, giving a second anticommuting operator that is in general nonunitary. This could be used directly in QSE, although the analysis of Sec. VI no longer holds unless $\hat{A}^\dagger \hat{A} \in \mathbb{P}^2$. For symmetry verification, we require the form of the projector $\hat{M}_a$ onto the correct $\hat{A}^\dagger \hat{A}|\psi\rangle = a|\psi\rangle$ subspace. This is a difficult task in general to construct (for $\hat{A}\hat{H}$, it is equivalent to diagonalizing the Hamiltonian). We have been unable to construct any further bounds on the performance of QSE as an error mitigation strategy for a general Hermitian operator, nor for an operator which neither commutes nor anticommutes with $\hat{H}$. This is, however, an interesting direction for future research.

-
- [1] D. Ristè, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira, and L. DiCarlo, *Nat. Commun.* **6**, 6983 (2015).
 - [2] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, *Nature (London)* **508**, 500 (2014).
 - [3] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, *Nature* **536**, 63 (2016).
 - [4] T. Monz, D. Nigg, E. A. Martinez, M. F. Brandl, P. Schindler, R. Rines, S. X. Wang, I. L. Chuang, and R. Blatt, *Science* **351**, 1068 (2016).
 - [5] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, *Nature (London)* **536**, 441 (2016).
 - [6] J. Preskill, *Quantum* **2**, 79 (2018).
 - [7] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Silipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, A. Kandala, A. Mezzacapo, P. Müller, W. Riess, G. Salis, J. Smolin, I. Tavernelli, and K. Temme, *Quantum Science and Technology* **3**, 030503 (2018).
 - [8] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, A. Fowler, B. Foxen, M. Giustina, R. Graff, E. Jeffrey, T. Huang, J. Kelly, P. Klimov, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis, *Science* **360**, 195 (2018).
 - [9] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. Love, A. Aspuru-Guzik, and J. O'Brien, *Nat. Commun.* **5**, 4213 (2014).
 - [10] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven, and Garnet Kin-Lic Chan, *Phys. Rev. X* **8**, 011044 (2018).
 - [11] D. Poulin, A. Kitaev, D. S. Steiger, M. B. Hastings, and M. Troyer, *Phys. Rev. Lett.* **121**, 010501 (2018).
 - [12] D. W. Berry, M. Kieferová, A. Scherer, Y. R. Sanders, G. H. Low, N. Wiebe, C. Gidney, and R. Babbush, *npj Quant. Inf.* **4**, 22 (2018).
 - [13] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, Garnet Kin-Lic Chan, and R. Babbush, *Phys. Rev. Lett.* **120**, 110501 (2018).
 - [14] K. Temme, S. Bravyi, and J. M. Gambetta, *Phys. Rev. Lett.* **119**, 180509 (2017).
 - [15] Y. Li and S. C. Benjamin, *Phys. Rev. X* **7**, 021050 (2017).
 - [16] A. Kandala, K. Temme, A. D. Corcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, *arXiv:1805.04492*.
 - [17] M. Otten and S. Gray, *arXiv:1806.07860*.
 - [18] S. Endo, S. C. Benjamin, and Y. Li, *Phys. Rev. X* **8**, 031027 (2018).
 - [19] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, *Phys. Rev. A* **95**, 042308 (2017).
 - [20] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, *Phys. Rev. X* **8**, 011021 (2018).
 - [21] D. Gottesman, *arXiv:0904.2557*.
 - [22] B. M. Terhal, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [23] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge Series on Information and the Natural Sciences (Cambridge University Press, Cambridge, 2000).
 - [24] A. Y. Kitaev, *arXiv:quant-ph/9511026*.
 - [25] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New J. Phys.* **18**, 023023 (2016).
 - [26] E. Farhi, J. Goldstone, and S. Gutmann, *arXiv:1411.4028*.
 - [27] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, *Phys. Rev. X* **6**, 031007 (2016).
 - [28] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature (London)* **549**, 242 (2017).
 - [29] R. M. Parrish, L. A. Burns, D. G. A. Smith, A. C. Simmonett, A. E. DePrince, E. G. Hohenstein, U. Bozkaya, A. Y. Sokolov, R. Di Remigio, R. M. Richard, J. F. Gonthier, A. M. James, H. R. McAlexander, A. Kumar, M. Saitow, X. Wang, B. P. Pritchard, P. Verma, H. F. Schaefer, K. Patkowski, R. A. King, E. F. Valeev, F. A. Evangelista, J. M. Turney, T. D. Crawford, and C. D. Sherrill, *J. Chem. Theory Comput.* **13**, 3185 (2017).

- [30] J. R. McClean, I. D. Kivlichan, K. J. Sung, D. S. Steiger, Y. Cao, C. Dai, E. Schuyler Fried, C. Gidney, B. Gimby, P. Gokhale, T. Häner, T. Hardikar, V. Havlíček, C. Huang, J. Izaac, Z. Jiang, X. Liu, M. Neeley, T. O'Brien, I. Ozfidan, M. D. Radin, J. Romero, N. Rubin, N. P. D. Sawaya, K. Setia, S. Sim, M. Steudtner, Q. Sun, W. Sun, F. Zhang, and R. Babbush, [arXiv:1710.07629](https://arxiv.org/abs/1710.07629).
- [31] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, *Mol. Phys.* **109**, 735 (2011).
- [32] D. Gottesman, *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, edited by S. P. Corney, R. Delbourgo, and P. D. Jarvis (International Press, Cambridge, MA, 1999), pp. 32–43.
- [33] The QUANTUMSIM density matrix simulator can be found at <https://github.com/quantumsim/>.
- [34] T. E. O'Brien, B. Tarasinski, and L. DiCarlo, *npj Quant. Inf.* **3**, 39 (2017).
- [35] S. Filipp, P. Maurer, P. J. Leek, M. Baur, R. Bianchetti, J. M. Fink, M. Göppl, L. Steffen, J. M. Gambetta, A. Blais, and A. Wallraff, *Phys. Rev. Lett.* **102**, 200402 (2009).
- [36] J. M. Chow, L. DiCarlo, J. M. Gambetta, A. Nunnenkamp, L. S. Bishop, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Phys. Rev. A* **81**, 062325 (2010).
- [37] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, *Quant. Inf. Comput.* **15**, 1 (2015).
- [38] I. M. Georgescu, S. Ashhab, and F. Nori, *Rev. Mod. Phys.* **86**, 153 (2014).
- [39] M. Suzuki, *Comm. Math. Phys.* **51**, 183 (1976).
- [40] The cluster operator for this system is a sum of eight four-qubit terms, however, the action of each term on the Hartree-Fock starting state is identical, so only one is needed.
- [41] S. McArdle, X. Yuan, and S. Benjamin, [arXiv:1807.02467](https://arxiv.org/abs/1807.02467).