



Universiteit
Leiden
The Netherlands

Optimization of quantum algorithms for near-term quantum computers

Bonet Monroig, X.

Citation

Bonet Monroig, X. (2022, November 2). *Optimization of quantum algorithms for near-term quantum computers*. *Casimir PhD Series*. Retrieved from <https://hdl.handle.net/1887/3485163>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3485163>

Note: To cite this publication please use the final published version (if applicable).

2. Low-cost error mitigation by symmetry verification

2.1. 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 [69–73], 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 [47, 74, 75]. In particular, for the area of digital quantum simulation, it has been suggested that variational quantum eigensolvers [39] may be sufficiently low-cost to be performed on ~ 50 qubits [76–79]. 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 non-scalable, 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 [49, 50, 80, 81], and on probabilistic error cancellation, where an ensemble of noisy circuits is applied such that they average to the target error-free circuit [50, 51]. 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 [82, 83].

In this chapter 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 [84, 85]. We develop multiple protocols to perform

2. Low-cost error mitigation by symmetry verification

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 [82]. 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.

2.2. 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. 2.2.1), 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}_S projects onto the target eigenspace \mathcal{S} . We note that phase estimation [86] 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 $\hat{M}_s|\psi\rangle = |\psi\rangle$, measuring $\{\hat{M}_i\}$ on the noisy state ρ and post-selecting will project to the state

$$\rho_s = \frac{\hat{M}_s \rho \hat{M}_s}{\text{Trace}[\hat{M}_s \rho]}. \quad (2.1)$$

Then, we have

$$\text{Trace}[\rho_s |\psi\rangle\langle\psi|] = \frac{\text{Trace}[\rho |\psi\rangle\langle\psi|]}{\text{Trace}[\hat{M}_s \rho]} \geq \text{Trace}[\rho |\psi\rangle\langle\psi|], \quad (2.2)$$

and our new state ρ_s has strictly greater overlap with the target $|\psi\rangle$ than the pre-selection ρ (unless $\hat{M}_s \rho \hat{M}_s = \rho$, 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 $\langle\psi(t)|$ 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 the an experiment.

2.2.1. 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 $\hat{S} \in \mathbb{F}^N$ in terms of its tensor factors; $\hat{S} = \otimes_i \hat{S}_i$, and let N_S be the number of non-trivial $\hat{S}_i = \{X, Y, Z\}$. 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. 2.1(a). For each non-trivial \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. 2.1(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 $O(N_S)$.

2. Low-cost error mitigation by symmetry verification

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. 2.2(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. 2.2(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. 2.2(a)) 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. 2.1(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.

2.2.2. 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 [39, 87].

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 non-superconducting 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 [39], and the quantum approximate optimization algorithm (QAOA) [88]. In the former, the ansatz is taken as




Figure 2.1.: 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).

2. Low-cost error mitigation by symmetry verification

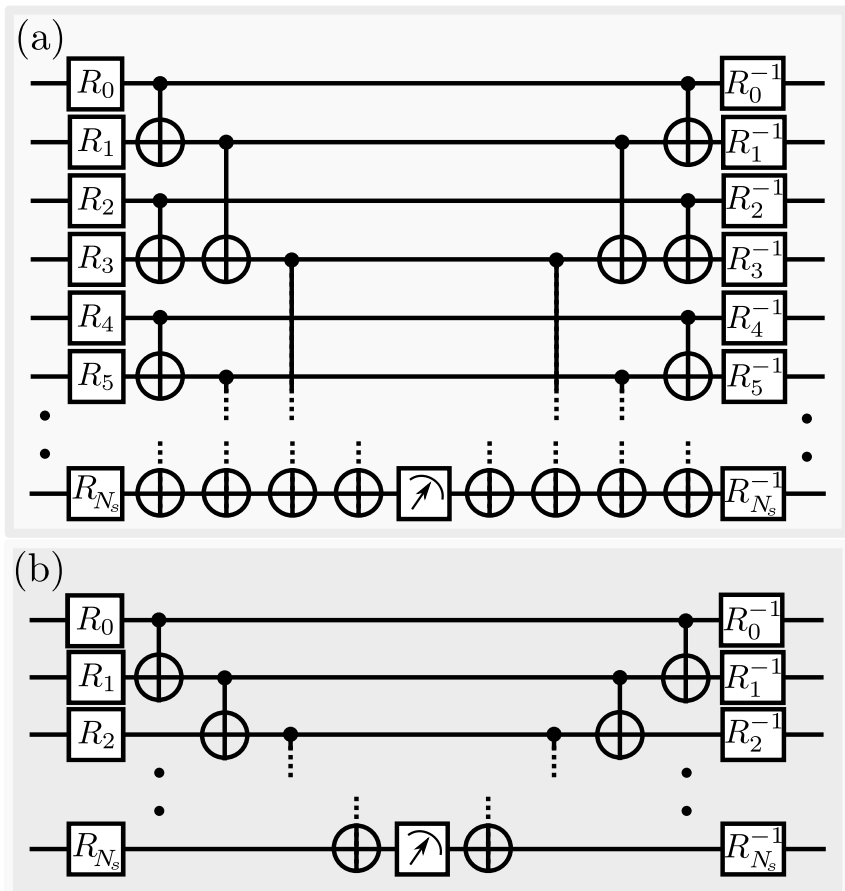


Figure 2.2.: Quantum circuits for in-line symmetry verification. (a) The optimal verification circuit has $O(\log(N_s))$ depth, but requires long-range connectivity between qubits, which is not available on many architectures. (b) In the presence of linear connectivity, an $O(N_s)$ depth verification circuit is optimal.

the expansion of the cluster operator $e^{\hat{T}-\hat{T}^\dagger}$

$$\hat{T} = \sum_n \hat{T}^{(n)}, \quad (2.3)$$

$$\hat{T}^{(n)} = \sum_{i_1, \dots, i_n; j_1, \dots, j_n} \theta_{j_1, \dots, j_n}^{i_1, \dots, i_n} \left(\prod_{m=1}^n \hat{c}_{i_m}^\dagger \right) \left(\prod_{m=1}^n \hat{c}_{j_m} \right) \quad (2.4)$$

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 semi-colon, and filled molecular orbitals to the right. This exponentiation is typically performed by the Trotter-Suzuki expansion, leaving a series of unitaries

$$\prod_{i;j} e^{\theta_j^i (\hat{c}_i^\dagger \hat{c}_j - \hat{c}_j^\dagger \hat{c}_i)} \prod_{i,j;k,l} e^{\theta_{k,l}^{i,j} (\hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_k \hat{c}_l - \hat{c}_l^\dagger \hat{c}_k^\dagger \hat{c}_j \hat{c}_i)} \dots \quad (2.5)$$

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

$$\rho_s = \hat{M}_s \rho \hat{M}_s, \quad \rho_r = (\mathbb{1} - \hat{M}_s) \rho (\mathbb{1} - \hat{M}_s). \quad (2.6)$$

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

$$\frac{\text{Trace}[\hat{H}\rho_r]}{\text{Trace}[\rho_r]} < \frac{\text{Trace}[\hat{H}\rho_s]}{\text{Trace}[\rho_s]}. \quad (2.7)$$

We can calculate

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

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

2. Low-cost error mitigation by symmetry verification

with high-energy states. As such, we would suggest that such a failure implies the energy of ρ itself is not to be trusted.

2.2.3. 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 $\hat{M}_s = \frac{1}{2}(1 + s\hat{S})$. Then, the expectation value of \hat{P} on the state ρ_s targeted by the symmetry verification can be expanded using Eq. 2.1

$$\begin{aligned} \text{Trace}[\hat{P}\rho_s] &= \text{Trace}\left[\hat{P} \frac{\hat{M}_s \rho \hat{M}_s}{\text{Trace}[\hat{M}_s \rho]}\right] \\ &= \frac{\text{Trace}[\hat{P}\rho] + s\text{Trace}[\hat{P}\hat{S}\rho]}{1 + s\text{Trace}[\hat{S}\rho]}, \end{aligned} \quad (2.8)$$

where we have used the cyclic property of the trace and the fact that $[\hat{P}, \hat{M}_s] = 0$ to write $\text{Trace}[\hat{P}\hat{M}_s \rho \hat{M}_s] = \text{Trace}[\hat{P}\hat{M}_s \rho]$, and expanded our definition of \hat{M}_s . The expectation values $\text{Trace}[\hat{S}\rho]$, $\text{Trace}[\hat{P}\rho]$, and $\text{Trace}[\hat{P}\hat{S}\rho]$ may be then calculated using the unverified state ρ , and substituted into Eq. 2.8 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 $\text{Trace}[\hat{P}\rho]$.

Post-selected symmetry verification can be observed to be identical to a form of the quantum subspace expansion (QSE) [82]. 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 $|\psi(\vec{\theta})\rangle$ to obtain a set of states $|\phi_i\rangle = \hat{E}_i|\psi(\vec{\theta})\rangle$. The spectrum of the Hamiltonian within the manifold spanned by these states can be calculated as the solution to the generalized eigenvalue problem

$$\hat{H}_{\text{QSE}}|\xi\rangle = \lambda \hat{B}_{\text{QSE}}|\xi\rangle. \quad (2.9)$$

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

$$[\hat{H}_{\text{QSE}}]_{i,j} = \text{Trace}[\hat{H}|\phi_i\rangle\langle\phi_j|], \quad (2.10)$$

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

$$[\hat{B}_{\text{QSE}}]_{i,j} = \text{Trace}[|\phi_i\rangle\langle\phi_j|], \quad (2.11)$$

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. 2.10 and 2.11 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 [82, 83]. However, if the set $\{\mathbb{1}, \hat{S}\}$ 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

$$\begin{aligned} \text{Trace}[\hat{H}\rho_s] &= \sum_i h_i \text{Trace}[\hat{P}_i\rho_s] \\ &= \sum_i \frac{\text{Trace}[h_i\hat{P}_i\rho] + s\text{Trace}[h_i\hat{P}_i\hat{S}\rho]}{1 + s\text{Trace}[\hat{S}\rho]} \\ &= \frac{\text{Trace}[\hat{H}\rho] + s\text{Trace}[\hat{H}\hat{S}\rho]}{1 + s\text{Trace}[\hat{S}\rho]}. \end{aligned} \quad (2.12)$$

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

$$\hat{H}_{\text{QSE}} = \begin{bmatrix} \text{Trace}[\hat{H}\rho] & \text{Trace}[\hat{H}\hat{S}\rho] \\ \text{Trace}[\hat{H}\hat{S}\rho] & \text{Trace}[\hat{H}\rho] \end{bmatrix}, \quad (2.13)$$

$$\hat{B}_{\text{QSE}} = \begin{bmatrix} 1 & \text{Trace}[\hat{S}\rho] \\ \text{Trace}[\hat{S}\rho] & 1 \end{bmatrix}. \quad (2.14)$$

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

$$\hat{B}_{\text{QSE}}^{-1}\hat{H}_{\text{QSE}} = \frac{1}{1 - \text{Trace}[\hat{S}\rho]^2} \begin{bmatrix} \alpha & \beta \\ \beta & \alpha \end{bmatrix}, \quad (2.15)$$

where

$$\alpha = \text{Trace}[\hat{H}\rho] - \text{Trace}[\hat{S}\rho]\text{Trace}[\hat{H}\hat{S}\rho], \quad (2.16)$$

$$\beta = \text{Trace}[\hat{H}\hat{S}\rho] - \text{Trace}[\hat{H}\rho]\text{Trace}[\hat{S}\rho]. \quad (2.17)$$

2. Low-cost error mitigation by symmetry verification

The eigenvalues of this matrix take the form

$$\lambda = \frac{1}{1 - \text{Trace}[\hat{S}\rho]^2}(\alpha \pm \beta) \quad (2.18)$$

$$= \frac{\text{Trace}[\hat{H}\rho] \pm \text{Trace}[\hat{H}\hat{S}\rho]}{1 \pm \text{Trace}[\hat{S}\rho]}, \quad (2.19)$$

which can be seen to be equal to those found in Eq. 2.12. 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 [82] to account for symmetries during QSE by projecting \hat{H}_{QSE} and \hat{B}_{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.

2.3. 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 [39, 41, 42, 83]. 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

$$\hat{H} = h_0\mathbb{1}\mathbb{1} + h_1\mathbb{1}Z + h_2Z\mathbb{1} + h_3XX + h_4YY + h_5ZZ, \quad (2.20)$$

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 [89] and OpenFermion [90] packages. The Hamiltonian can be seen to commute with the symmetry $\hat{S} = ZZ$. Our ground state wavefunction has

2.3. Simulation of symmetry verification on the hydrogen molecule

non-trivial overlap with the Hartree-Fock wavefunction, which is in the $ZZ = -1$ subspace; this is then our target subspace. We follow the unitary coupled cluster ansatz of [41], which consists of exciting our system to the $|01\rangle$ state, and performing the unitary rotation

$$\hat{U}(\theta) = e^{-i\theta X_0 Y_1}. \quad (2.21)$$

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 [39].

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. 2.1(a). The in-line symmetry verification is performed in a manner similar to Fig. 2.2(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 $\text{Trace}[\rho \hat{P}]$, we can propagate \hat{P} through the symmetry verification circuit [91] and measure the corresponding Pauli term. It is then sufficient to rotate the control qubit to recover the expectation values $\langle \mathbb{1}Z \rangle$ and $\langle XX \rangle$. From this we may calculate all other expectation values in Eq. 2.20 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 `quantumsim` density matrix simulator [92]. We take gate error models and parameters similar to previous simulation work based on experimental data of state-of-the-art superconducting transmon qubits [92]. 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 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 cancelled by linear inversion tomography [93, 94].

Using the above error model, we observe (Fig. 2.3) that the un-mitigated VQE (blue points) achieves an error in the energy of approximately 0.01 –

2. Low-cost error mitigation by symmetry verification

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 five-fold improvement over the unmitigated circuit, while ancilla (orange crosses) and in-line (green squares) symmetry verification show an approximately two-fold and three-fold 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 [84]. 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. 2.4 we calculate the error in determining the ground state energy near the minima of the bond dissociation curve (0.75Å 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 $T_\phi = 20 \mu\text{s}$ ($T_1 = 20 \mu\text{s}$) fixed. In the absence of error 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.

2.4. 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




Figure 2.3.: (Color online) Accuracy of the VQE over the entire bond dissociation curve using the different symmetry verification methods mentioned in the text (labelled 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.

2. Low-cost error mitigation by symmetry verification




Figure 2.4.: (Color online) Effect of varying decoherence times on the VQE accuracy. With all other error sources turned off, T_1 is varied with $T_\phi = 20 \mu\text{s}$ fixed (red-dashed curves), and T_ϕ is varied with $T_1 = 20 \mu\text{s}$ 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 $T_1 = T_\phi = 20 \mu\text{s}$, as can be seen from the complete overlap.

2.4. Inserting and rotating symmetries

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}_{ext}

$$\hat{H}_{\text{ext}} = \begin{bmatrix} \hat{H} & 0 \\ 0 & \hat{P}\hat{H}\hat{P} \end{bmatrix}. \quad (2.22)$$

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

$$\begin{bmatrix} 0 & \hat{P} \\ \hat{P} & 0 \end{bmatrix} = X\hat{P}, \quad (2.23)$$

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

$$\hat{Q}_{\text{ext}} = \begin{cases} \mathbb{I}\hat{Q} & \text{if } [\hat{Q}, \hat{P}] = 0 \\ X\hat{Q} & \text{if } \{\hat{Q}, \hat{P}\} = 0 \end{cases}. \quad (2.24)$$

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

$$\prod_j e^{i\theta_j \hat{Q}_j}, \quad \hat{Q}_j \in \mathbb{P}^N, \quad (2.25)$$

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}_{ext} (as per Eq. 2.24), and re-decomposing the operations into a circuit (using e.g. the methods of [31, 95]). 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. 2.24), that commute with both \hat{H}_{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. 2.25. This is the case for traditional Hamiltonian simulation [96], quantum phase estimation [31], and the UCC QSE discussed previously, all of which require exponentiating an operator via the Suzuki-Trotter expansion [97].

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

2. Low-cost error mitigation by symmetry verification

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. $\mathbb{1}Z, Z\mathbb{1}$, and ZZ are equivalent to XX, YY and ZZ , but not to $\mathbb{1}X, \mathbb{1}Y$ and $\mathbb{1}Z$). 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

$$\hat{P} \in \mathbb{P} \rightarrow \hat{R}^\dagger \hat{P} \hat{R} = \begin{cases} \hat{P} & \text{if } [\hat{P}, \hat{Q}] = 0 \\ i\hat{P}\hat{Q} & \text{if } \{\hat{P}, \hat{Q}\} = 0 \end{cases}. \quad (2.26)$$

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.

2.5. 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

$$\begin{aligned} \hat{H} = & h_I \mathbb{1} + \sum_i h_i Z_i + \frac{1}{2} \sum_{i \neq j} h_{i,j} Z_i Z_j \\ & + h_s (X_0 Y_1 Y_2 X_3 + Y_0 X_1 X_2 Y_3 \\ & - X_0 X_1 Y_2 Y_3 - Y_0 Y_1 X_2 X_3), \end{aligned} \quad (2.27)$$

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 choose again the unitary coupled cluster ansatz for the VQE, which can be reduced to the operator¹

$$\hat{U}(\theta) = e^{i\theta Y_0 X_1 X_2 X_3}. \quad (2.28)$$

¹The cluster operator for this system is a sum of 8 four-qubit terms, however the action of each term on the Hartree-Fock starting state is identical, so only one is needed.

2.5. Extending the symmetry verification of the hydrogen molecule

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

$$\hat{R} = e^{i\frac{\pi}{2}Y_0X_2}e^{i\frac{\pi}{2}Y_1X_3}. \quad (2.29)$$

This transforms the symmetry $\hat{S}_0 \rightarrow X_0X_1X_2X_3$, whilst 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

$$\hat{U}(\theta) \rightarrow \hat{R}\hat{U}\hat{R}^\dagger = e^{i\theta Y_0Z_1X_2}. \quad (2.30)$$

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. 2.5, we compare the performance of the two different circuits above to the two-qubit circuit of Fig. 2.3, with and without the addition of S-QSE. At small bond distance ($\lesssim 0.75\text{\AA}$), 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 ($\gtrsim 0.75\text{\AA}$), the ground state is prone to T_2 noise, at which point we see the rotated 4-qubit S-QSE simulation significantly outperforming its counterparts. At the largest distance studied, this simulation achieves a two-fold 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.

2. Low-cost error mitigation by symmetry verification




Figure 2.5.: (Color online) Adding and adjusting symmetries to optimize symmetry verification. The blue (dots) and red (diamonds) curves correspond to their coloured (shaped) counterparts in Fig. 2.3, whilst 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).

2.6. Conclusion

In this paper we have presented a new 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 mid-range 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 maximise 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 mid-size NISQ circuits and large-scale fault-tolerant ones.

Appendix

2.A. 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 anti-commutes 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

$$|\pm\rangle = \frac{1}{\sqrt{2}}(|\psi\rangle \pm \hat{A}|\psi\rangle). \quad (2.31)$$

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

$$\hat{B}_{\text{QSE}} = \begin{bmatrix} 1 & \text{Trace}[\hat{A}\rho] \\ \text{Trace}[\hat{A}^\dagger\rho] & 1 \end{bmatrix}. \quad (2.32)$$

$$\hat{H}_{\text{QSE}} = \begin{bmatrix} \text{Trace}[\hat{H}\rho] & \text{Trace}[\hat{H}\hat{A}\rho] \\ \text{Trace}[-\hat{H}\hat{A}^\dagger\rho] & -\text{Trace}[\hat{H}\rho] \end{bmatrix}. \quad (2.33)$$

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

$$\hat{B}_{\text{QSE}}^{-1}\hat{H}_{\text{QSE}} = \frac{1}{1 - |\text{Trace}[\hat{A}\rho]|^2} \begin{bmatrix} \alpha & \beta \\ -\beta^* & -\alpha^* \end{bmatrix}, \quad (2.34)$$

where

$$\alpha = \text{Trace}[\hat{H}\rho] + \text{Trace}[\hat{A}\rho]\text{Trace}[\hat{H}\hat{A}\rho] \quad (2.35)$$

$$\beta = \text{Trace}[\hat{H}\hat{A}\rho] + \text{Trace}[\hat{H}\rho]\text{Trace}[\hat{A}\rho]. \quad (2.36)$$

2. Low-cost error mitigation by symmetry verification

The solution to the equation is

$$E_{\text{QSE}}^2 = \frac{|\alpha|^2 + |\beta|^2}{(1 - |\text{Trace}[\hat{A}\rho]|^2)^2} \quad (2.37)$$

$$= \frac{\text{Trace}[\hat{H}\rho]^2 + |\text{Trace}[\hat{H}\hat{A}\rho]|^2}{1 - |\text{Trace}[\hat{A}\rho]|^2}. \quad (2.38)$$

To understand the gain in energy, $|\text{Trace}[\hat{H}\hat{A}\rho]|^2$, 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,

$$\rho = |a|^2|\psi\rangle\langle\psi| + |b|^2\hat{A}|\psi\rangle\langle\psi|\hat{A}, \quad (2.39)$$

$\text{Trace}[\hat{H}\hat{A}\rho] = \text{Trace}[\hat{A}\rho] = 0$ (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:

$$\begin{aligned} \rho = & (\cos(\theta)|\psi\rangle + \sin(\theta)e^{i\phi}\hat{A}|\psi\rangle) \\ & \times (\cos(\theta)\langle\psi| + \sin(\theta)e^{-i\phi}\langle\psi|\hat{A}^\dagger). \end{aligned} \quad (2.40)$$

We can calculate

$$\text{Trace}[\hat{H}\rho] = E \cos(2\theta), \quad (2.41)$$

$$\text{Trace}[\hat{A}\rho] = \sin(2\theta)(1 + Ae^{i\phi}), \quad (2.42)$$

$$\text{Trace}[\hat{H}\hat{A}\rho] = E \sin(2\theta)(Ae^{i\phi} - 1), \quad (2.43)$$

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

$$E_{\text{QSE}}^2 = E^2 \frac{\cos^2(2\theta) + \sin^2(2\theta)\chi_+}{1 - \sin^2(2\theta)\chi_-}, \quad (2.44)$$

$$\chi_{\pm} = (1 \pm Ae^{i\phi})(1 \pm Ae^{-i\phi}). \quad (2.45)$$

We see that if $A = 1, \phi = \frac{\pi}{2}$, QSE corrects the coherent error entirely, whilst 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 [82] 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 $\{\hat{A}, \hat{H}\} = 0$, $\{\hat{A}\hat{H}, \hat{H}\} = 0$ as well, giving a second anticommuting operator that is in general non-unitary. This could be

2.A. Appendix: Error mitigation of QSE with anticommuting operators

used directly in QSE, although the analysis of Sec. 2.2.3 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 anti-commutes with \hat{H} . This is, however, an interesting direction for future research.

