

Unitary partitioning and the contextual subspace variational quantum eigensolver

Alexis Ralli  ^{1,*} Tim Weaving  ^{1,†} Andrew Tranter, ^{2,3,‡} William M. Kirby, ^{2,§} Peter J. Love, ^{2,4,||} and Peter V. Coveney ^{1,5,¶}

¹Centre for Computational Science, Department of Chemistry, University College London, London WC1H 0AJ, United Kingdom

²Department of Physics and Astronomy, Tufts University, Medford, Massachusetts 02155, USA

³Cambridge Quantum Computing, 9a Bridge Street, Cambridge CB2 1UB, United Kingdom

⁴Computational Science Initiative, Brookhaven National Laboratory, Upton, New York 11973, USA

⁵Informatics Institute, University of Amsterdam, Amsterdam 1098 XH, Netherlands



(Received 13 August 2022; accepted 18 January 2023; published 9 February 2023)

The contextual subspace variational quantum eigensolver (CS-VQE) is a hybrid quantum-classical algorithm that approximates the ground-state energy of a given qubit Hamiltonian. It achieves this by separating the Hamiltonian into contextual and noncontextual parts. The ground-state energy is approximated by classically solving the noncontextual problem, followed by solving the contextual problem using VQE, constrained by the noncontextual solution. In general, computation of the contextual correction needs fewer qubits and measurements compared with solving the full Hamiltonian via traditional VQE. We simulate CS-VQE on different tapered molecular Hamiltonians and apply the unitary partitioning measurement reduction strategy to further reduce the number of measurements required to obtain the contextual correction. Our results indicate that CS-VQE combined with measurement reduction is a promising approach to allow feasible eigenvalue computations on noisy intermediate-scale quantum devices. We also provide a modification to the CS-VQE algorithm; the CS-VQE algorithm previously could cause an exponential increase in Hamiltonian terms but with this modification now at worst will scale quadratically.

DOI: 10.1103/PhysRevResearch.5.013095

I. INTRODUCTION

One of the fundamental goals of quantum chemistry is to solve the time-independent nonrelativistic Schrödinger equation. The eigenvalues and eigenvectors obtained allow different molecular properties to be studied from first principles. Standard methods project the problem onto a Fock space (η electrons distributed in M orbitals) and solve. Under this approximation, the problem scales exponentially with system size, where the number of Slater determinants (configurations) scales as $\binom{M}{\eta}$ making the problem classically intractable [1]. Quantum computers can efficiently represent the full configuration interaction (FCI) Hilbert space and offer a potential way to efficiently solve such molecular problems [2,3]. This use case is often the canonical example of where the first quantum computers will be advantageous over conventional computers [4,5].

In the fault-tolerant regime, quantum phase estimation (QPE) [6] provides a practical way to perform quantum chemistry simulations in polynomial time [2]. However, current noisy intermediate-scale quantum (NISQ) devices cannot implement this algorithm due to the deep quantum circuits and long coherence times required [7,8].

The constraints on present-day devices have given rise to a family of quantum-classical algorithms that leverage as much classical processing as possible to reduce the quantum resources required to solve the problem at hand. Common examples of NISQ algorithms are the variational quantum eigensolver (VQE) [9], quantum approximate optimization algorithm (QAOA) [10], and variational quantum linear solver (VQLS) [11]. A good example is the recently proposed entanglement forging method [12], where the electronic structure problem for H_2O was reduced from a 10-qubit problem to multiple 5-qubit problems that were each studied using conventional VQE and classically combined. Recently, another novel approach known as the quantum-classical hybrid quantum Monte Carlo (QC-QMC) method was used to unbiased the sign problem in the projector Monte Carlo (PMC) method, which implements imaginary time evolution [13]. At a high level, the accuracy of a constrained PMC calculation is determined by the quality of trial wave functions. Quantum computers offer a way to efficiently store highly entangled trial wave functions and measure certain overlaps, which would require exponential resources classically. Huggins *et al.* performed QC-QMC simulations of different chemical systems on Google's Sycamore processor and obtained results competitive with state-of-the-art classical methods [13].

*alexis.ralli.18@ucl.ac.uk

†timothy.weaving.20@ucl.ac.uk

‡tufts@atranter.net

§william.kirby@tufts.edu

||peter.love@tufts.edu

¶p.v.coveney@ucl.ac.uk

The contextual subspace VQE (CS-VQE) algorithm is another hybrid quantum-classical approach [14]. It gives an approximate simulation method, where the quantum resources required can be varied for a trade-off in accuracy. This allows problems to be studied where the full Hamiltonian would normally be too large to investigate on current NISQ hardware. This was shown in the original CS-VQE paper, where chemical accuracy for various molecular systems was reached using significantly fewer qubits compared with the number required for VQE on the full system [14]. As CS-VQE reduces the number of qubits required for simulation, the number of terms in a Hamiltonian requiring separate measurements is also reduced.

A natural question that arises from this is whether measurement reduction schemes can be utilized to reduce the overall measurement cost of these already reduced CS-VQE Hamiltonians [15–28]. The goal of this work was to investigate the possible reductions given by the unitary partitioning strategy [15,29,30] and whether chemical accuracy on larger molecules can be reached on currently available NISQ hardware.

This paper is structured as follows. Section II summarizes the CS-VQE algorithm. Here we provide a modification to the unitary partitioning step of the CS-VQE algorithm; the CS-VQE algorithm previously could cause the number of terms in a Hamiltonian to exponentially increase but with this modification now will at worst cause a quadratic increase. Section III is split into a description of the method in Sec. III A and two main parts, Secs. III B and III C. Section III B examines a model problem to exemplify each step of the CS-VQE algorithm. Section III C gives the numerical results of applying unitary partitioning measurement reduction to a test bed of different molecular structure Hamiltonians, where the contextual subspace approximation has been employed.

II. BACKGROUND

To keep our discussion self-contained and establish notation, we summarize the necessary background theory of the contextual subspace VQE algorithm in this section.

A. Contextuality

The foundation of quantum contextuality is the Bell-Kochen-Specker (BKS) theorem [31]. In lay terms, every measurement provides a classical probability distribution (via the spectral theorem), and a joint distribution can be built as a product over all possible measurements [32]. The BKS theorem proves that it is impossible to reproduce the probabilities of every possible measurement outcome for a quantum system as marginals of this joint probability distribution [33]. This is related to how quantum mechanics does not allow models that are locally causal in a classical sense [34]. Contextuality is a generalization of nonlocality [34,35]. This means that quantum measurement cannot be understood as simply revealing a preexisting value of some underlying hidden variable [36,37]. Bell's theorem also reaches a similar conclusion against hidden variables [38], but in a different way.

A good example of this phenomenon is the “Peres-Mermin square” [37,39], where no state preparation is involved and

only observables are considered. We include an example in the Appendix and remark on the relation to VQE. Colloquially, for a noncontextual problem it is possible to assign deterministic outcomes to observables simultaneously without contradiction; however, for a contextual problem this is not possible [14].

The following sections set out the contextual subspace VQE algorithm, and we provide an alternate way to construct $U_{\mathcal{W}}$ (defined below) compared with the original work [14]. This modification addresses the exponential scaling part of the method. Further background on the full CS-VQE algorithm is given in the Supplemental Material [40].

B. Contextual subspace VQE

Consider a Hamiltonian expressed as

$$H_{\text{full}} = \sum_a c_a P_a = \sum_a c_a \left(\bigotimes_{j=0}^{n-1} \sigma_j^{(a)} \right) = \sum_a c_a (\sigma_0^{(a)} \otimes \sigma_1^{(a)} \otimes \cdots \otimes \sigma_{n-1}^{(a)}), \quad (1)$$

where c_a are real coefficients. Each Pauli operator P_a is made up of an n -fold tensor product of single-qubit Pauli matrices $\sigma_j \in \{\mathcal{I}, X, Y, Z\}$, where j indexes the qubit the operator acts on. The CS-VQE algorithm is based on separating such a Hamiltonian into a contextual part and a noncontextual part [14]:

$$H_{\text{full}} = H_{\text{con}} + H_{\text{noncon}}. \quad (2)$$

As it is possible to assign definite values to all terms in H_{noncon} without contradiction, a classical hidden-variable model (or quasiquantized model) can be used to represent this system [41].

In Ref. [42], such a model is constructed along with a classical algorithm to solve it. This was based on the work of Spekkens [43,44]. Solving this model yields a noncontextual ground state.

Once that solution is obtained, the remaining contextual part of the problem is solved. Solutions to H_{con} must be consistent with the noncontextual ground state, which defines a subspace of allowed states [14]. By projecting the problem into this subspace the overall energy is given by

$$\begin{aligned} E(\vec{\theta}; \vec{q}, \vec{r}) &= E_{\text{noncon}}(\vec{q}, \vec{r}) + E_{\text{con}}(\vec{\theta}; \vec{q}, \vec{r}) \\ &= E_{\text{noncon}}(\vec{q}, \vec{r}) \\ &\quad + \frac{\langle \psi_{\text{con}}(\vec{\theta}) | Q_{\mathcal{W}}^{\dagger} U_{\mathcal{W}}^{\dagger} H_{\text{con}} U_{\mathcal{W}} Q_{\mathcal{W}} | \psi_{\text{con}}(\vec{\theta}) \rangle}{\langle \psi_{\text{con}}(\vec{\theta}) | Q_{\mathcal{W}}^{\dagger} Q_{\mathcal{W}} | \psi_{\text{con}}(\vec{\theta}) \rangle} \\ &= \frac{\langle \psi_{\text{con}}(\vec{\theta}) | Q_{\mathcal{W}}^{\dagger} U_{\mathcal{W}}^{\dagger} H_{\text{full}} U_{\mathcal{W}} Q_{\mathcal{W}} | \psi_{\text{con}}(\vec{\theta}) \rangle}{\langle \psi_{\text{con}}(\vec{\theta}) | Q_{\mathcal{W}}^{\dagger} Q_{\mathcal{W}} | \psi_{\text{con}}(\vec{\theta}) \rangle}. \end{aligned} \quad (3)$$

We have written $U_{\mathcal{W}}$ rather than $U_{\mathcal{W}}(\vec{q}, \vec{r})$ to simplify our notation.

The vector (\vec{q}, \vec{r}) should be thought of as parameters that define a particular noncontextual state: Normally, this will be a parametrization for the noncontextual ground state [14,42].

This vector has a size of at most $2n + 1$ for a Hamiltonian defined on n qubits [42]. From (\vec{q}, \vec{r}) , we define a set of stabilizers \mathcal{W} which stabilize that particular noncontextual state [14]. The unitary $U_{\mathcal{W}}$ maps each of these stabilizers to a distinct single-qubit Pauli matrix; details of this are covered in Sec. II E. By enforcing the eigenvalue of these single-qubit Pauli operators we define a subspace of allowed quantum states that are consistent with the noncontextual state. To constrain the problem to this subspace, we use the projector $Q_{\mathcal{W}}$. Note that this is not a unitary operation: hence the renormalization in Eq. (3). By projecting our contextual Hamiltonian into this subspace, $H_{\text{con}} \mapsto H_{\text{con}}^{\mathcal{W}} = Q_{\mathcal{W}}^{\dagger} U_{\mathcal{W}}^{\dagger} H_{\text{con}} U_{\mathcal{W}} Q_{\mathcal{W}}$, we ensure that solutions to $H_{\text{con}}^{\mathcal{W}}$ remain in the subspace consistent with the noncontextual solution [14]. In other words, this operation means that solutions to $H_{\text{con}}^{\mathcal{W}}$ will remain consistent with the noncontextual solution. Section II E goes into detail on this.

The contextual trial or ansatz state is prepared as $|\psi_{\text{con}}(\vec{\theta})\rangle = U_{\mathcal{W}}^{\dagger} V(\vec{\theta}) |0\rangle^{\otimes n}$, where $V(\vec{\theta})$ is the parametrized operator that prepares it. The projector $Q_{\mathcal{W}}$, in Eq. (3), then projects this state into the subspace of possible states consistent with the noncontextual ground state. Again this depends on which stabilizer eigenvalues are fixed. Note that as $Q_{\mathcal{W}}$ is not a unitary operation, the state must be renormalized. Further analysis of the contextual subspace VQE projection ansatz is provided in Ref. [45]. In Sec. II C we discuss how to solve the noncontextual problem.

C. Noncontextual Hamiltonian

For a given noncontextual Hamiltonian, we define $\mathcal{S}^{H_{\text{noncon}}}$ to be the set of P_i present in H_{noncon} . This set can be expanded as two subsets denoted as \mathcal{Z} and \mathcal{T} , representing the set of fully commuting operators and its complement, respectively [14,42]. The set \mathcal{T} can be expanded into N cliques, where operators within a clique must all commute with each other and operators between cliques must pairwise anticommute. This is because commutation forms an equivalence relation on \mathcal{T} if and only if $\mathcal{S}^{H_{\text{noncon}}}$ is noncontextual [14,42].

A hidden-variable model for such a system can be built, where the set of observables \mathcal{R} that define the phase-space points of the hidden-variable model is [14,42,46]

$$\begin{aligned} \mathcal{R} &\equiv \{P_0^{(j)} | j = 0, 1, \dots, N-1\} \cup \{G_0, G_1, G_2, \dots\} \\ &\equiv \{P_0^{(j)} | j = 0, 1, \dots, N-1\} \cup \mathcal{G}. \end{aligned} \quad (4)$$

The set \mathcal{G} represents an independent set of Pauli operators that generates the set of commuting observables \mathcal{Z} . Each $P_0^{(j)}$ corresponds to a chosen Pauli operator in the j th clique of \mathcal{T} : By convention we say that this is the first operator in the set, but this can be any operator in the j th clique.

With respect to the phase-space model given in Ref. [42], a valid noncontextual state is defined by the parameters (\vec{q}, \vec{r}) , which set the expectation value of the operators in \mathcal{R} [Eq. (4)]. Each operator in \mathcal{G} is assigned the value $\langle G_i \rangle = q_i = \pm 1$. The operators in \mathcal{T} are assigned the values $\langle P_0^{(j)} \rangle = r_j$, where \vec{r} is a unit vector ($|\vec{r}| = 1$) [14,42]. The number of elements in \vec{q} and \vec{r} are $|\mathcal{G}|$ and N , respectively. For n qubits the size of $|\mathcal{R}|$ is bounded by $2n + 1$, which bounds the size of (\vec{q}, \vec{r}) [42,46].

The observables for the N anticommuting $P_0^{(j)}$ operators in \mathcal{R} can be combined into the observable [42]

$$A(\vec{r}) = \sum_{j=0}^{N-1} r_j P_0^{(j)}. \quad (5)$$

We denote the set of Pauli operators making up this operator as $\mathcal{A} \equiv \{P_0^{(j)} | j = 0, 1, \dots, N-1\}$.

The expectation value of $A(\vec{r})$ is assigned by the hidden-variable model to always be $+1$, due to

$$\langle A(\vec{r}) \rangle = \sum_{j=0}^{N-1} r_j \langle P_0^{(j)} \rangle = \sum_{j=0}^{N-1} r_j r_j = \sum_{j=0}^{N-1} |r_j|^2 = +1, \quad (6)$$

using $\langle P_0^{(j)} \rangle = r_j$ and $|\vec{r}| = 1$.

The expectation value for H_{noncon} can be induced, by setting the expectation values of operators in \mathcal{R} [Eq. (4)], as this set generates $\mathcal{S}^{H_{\text{noncon}}}$ [14,42]. To find the ground state of H_{noncon} , we perform a brute-force search over this space. For each possible ± 1 combination of expectation values for each G_j ($2^{|\mathcal{G}|}$ possibilities), the energy is minimized with respect to the unit vector \vec{r} : That sets the expectation values $\langle P_0^{(j)} \rangle = r_j$. The Supplemental Material provides further algorithmic details [40]. The vector (\vec{q}, \vec{r}) that was found to give the lowest energy defines the noncontextual ground state. Each noncontextual state (\vec{q}, \vec{r}) corresponds to subspaces of quantum states, which we will describe in Sec. II D.

D. Contextual subspace

In Refs. [15,29] it was shown that an operator constructed as a normalized linear combination of pairwise anticommuting Pauli operators, such as $A(\vec{r})$ [Eq. (5)], is equivalent to a single Pauli operator up to a unitary rotation R . We can therefore write $A(\vec{r}) \mapsto P_0^{(k)} = RA(\vec{r})R^{\dagger}$ for a selected $P_0^{(k)} \in \mathcal{A}$. We write the set $\mathcal{R}_{\vec{r}}$ [Eq. (4)] under this transformation:

$$\mathcal{R}_{\vec{r}} \equiv \{A(\vec{r})\} \cup \mathcal{G} \mapsto \mathcal{R}' \equiv \underbrace{\{P_0^{(k)}\}}_{P_0^{(k)} = RA(\vec{r})R^{\dagger}} \cup \mathcal{G}. \quad (7)$$

It will be shown later that the unitary R is constructed from the operators in \mathcal{A} . This means that the terms in \mathcal{G} are unaffected by this transformation, as operators in \mathcal{G} must universally commute and so must commute with R .

A given noncontextual state (\vec{q}, \vec{r}) is equivalent to the joint expectation value assignment of $\langle G_i \rangle = q_i = \pm 1$ and $\langle A(\vec{r}) \rangle = +1$. This defines a set of stabilizers:

$$\mathcal{W}_{\text{all}} \equiv \{q_0 G_0, q_1 G_1, \dots, q_{|\mathcal{G}|-1} G_{|\mathcal{G}|-1}, A(\vec{r})\}, \quad (8)$$

which by definition must stabilize that noncontextual state (\vec{q}, \vec{r}) or more precisely, the subspace of quantum states corresponding to it [47]. Note that $A(\vec{r})$ is not a conventional stabilizer, but is unitarily equivalent to a single-qubit operator $P_0^{(k)}$ [15,29].

We can consider this problem under the unitary transform defined in Eq. (7). The stabilizers in \mathcal{W}_{all} become

$$\mathcal{W}'_{\text{all}} \equiv \{q_0 G_0, q_1 G_1, \dots, q_{|\mathcal{G}|-1} G_{|\mathcal{G}|-1}, \xi P_0^{(k)}\}, \quad (9)$$

which defines a regular set of stabilizers for the noncontextual state (\vec{q}, \vec{r}) , which defines a subspace of quantum states. Here,

$\xi = \pm 1$, is determined by $RA(\vec{r})R^\dagger = \xi P_0^{(k)}$, and can always be chosen to be $+1$, which we do throughout this paper.

Altogether, when certain noncontextual stabilizers are fixed (by the noncontextual state), they specify a subspace of allowed quantum states that will be consistent with that noncontextual state and thus define the constraints for the contextual part of the problem. We refer to this subspace as the contextual subspace [14].

E. Mapping a contextual subspace to a stabilizer subspace

In CS-VQE, the expectation value of the full Hamiltonian is obtained according to Eq. (3). First, the noncontextual problem is solved yielding the noncontextual state (\vec{q}, \vec{r}) : normally the ground state (\vec{q}_0, \vec{r}_0) . The Supplemental Material shows how (\vec{q}_0, \vec{r}_0) can be obtained via a brute-force approach [40]. Next the contextual Hamiltonian is projected into the subspace of allowed quantum states consistent with the defined noncontextual state. This constraint is imposed via $H_{\text{full}} \mapsto H_{\text{full}}^{\mathcal{W}} = Q_{\mathcal{W}}^\dagger U_{\mathcal{W}}^\dagger H_{\text{full}} U_{\mathcal{W}} Q_{\mathcal{W}}$, where the expectation value is then found on a quantum device.

The unitary operation $U_{\mathcal{W}}$ is defined by the set of contextual stabilizers $\mathcal{W} \subseteq \mathcal{W}_{\text{all}}$ [Eq. (8)], whose eigenvalue we fix according to the noncontextual state. If $A(\vec{r}) \in \mathcal{W}$, meaning that $\langle A(\vec{r}) \rangle$ is fixed to be $+1$, then the steps summarized in Eq. (7) must first be performed to reduce $A(\vec{r})$ to a single Pauli operator. Clifford operators $V_i(P)$ are then used to map each $P \in \mathcal{W}$ to a single-qubit Z operator. Each V_i is made up of at most two $\frac{\pi}{2}$ Clifford rotations, generated by Pauli operators, per element in \mathcal{W} . In Ref. [14] it was shown that at most there will be $2n$ of these rotations, where n is the number of qubits the problem is defined on [48]. We can write this operator as

$$U_{\mathcal{W}}^\dagger(\vec{q}, \vec{r}) = \begin{cases} \prod_{P_i \in \mathcal{W} \subseteq \mathcal{W}_{\text{all}}} V_i(P_i) & \text{if } A(\vec{r}) \notin \mathcal{W} \\ (\prod_{P_i \in \mathcal{W} \subseteq \mathcal{W}_{\text{all}}} V_i(P_i))R & \text{if } A(\vec{r}) \in \mathcal{W}. \end{cases} \quad (10)$$

Applying $U_{\mathcal{W}}^\dagger \mathcal{W} U_{\mathcal{W}} = \mathcal{W}^Z$ results in a set of single-qubit Z Pauli operators. An implementation note is that each operator $V_i(P_i)$ in $U_{\mathcal{W}}$ depends on the others: This can be seen by expanding $U_{\mathcal{W}}^\dagger \mathcal{W} U_{\mathcal{W}}$. Therefore each V_i operator is dependent on the stabilizers in \mathcal{W} and the order in which they occur. We recursively define each V_i as follows.

- (1) Set $\mathcal{W} = RWR^\dagger$ if and only if $A(\vec{r}) \in \mathcal{W}$.
- (2) Find the unitary V_0 mapping the first Pauli operator $P_0 \in \mathcal{W}$ to a single-qubit Pauli operator.
- (3) Apply this operator to each operator in the set: $V_0 \mathcal{W} V_0^\dagger = \mathcal{W}^{(0)}$.
- (4) Find the unitary V_1 mapping $V_0 P_1 V_0^\dagger \in \mathcal{W}^{(0)}$ to a single-qubit Z Pauli operator.
- (5) Apply this operator to all operators in the set: $V_1 \mathcal{W}^{(0)} V_1^\dagger = \mathcal{W}^{(1)}$.
- (6) Repeat this procedure from step (3) until all the operators are mapped to single-qubit Z Pauli operators: $\mathcal{W} \mapsto \mathcal{W}^Z$.

Finally, the eigenvalue of each single-qubit Z Pauli stabilizers in \mathcal{W}^Z is defined by the vector \vec{q} of the noncontextual ground state (\vec{q}, \vec{r}) ; note that $\langle A(\vec{r}) \rangle$ is fixed to $+1$ and thus \vec{r} is not important here. $U_{\mathcal{W}}$ can flip the sign of these assignments, but it is efficient to classically determine by tracking how $U_{\mathcal{W}}$ affects the sign of the operators in \mathcal{W} .

To project the Hamiltonian into the subspace consistent with the noncontextual state, we first perform the following rotation: $H_{\text{full}} \mapsto H'_{\text{full}} = U_{\mathcal{W}}^\dagger H_{\text{full}} U_{\mathcal{W}}$. As this is a unitary transform, the resultant operator has the same spectrum as before. We then restrict the rotated Hamiltonian to the correct subspace by enforcing the eigenvalue of the operators in \mathcal{W}^Z , where the outcomes are defined by the noncontextual state. As each operator in \mathcal{W}^Z only acts nontrivially on a unique qubit, each stabilizer fixes the state of that qubit to be either $|0\rangle$ or $|1\rangle$. We write this state as

$$|\psi_{\text{fixed}}\rangle = \bigotimes_{P_v \in \mathcal{W}^Z} |i\rangle_v \begin{cases} i = 0 & \text{if } \langle P_v \rangle = +1 \\ i = 1 & \text{if } \langle P_v \rangle = -1, \end{cases} \quad (11)$$

where v indexes the qubit a given single-qubit stabilizer acts on and $\langle P_v \rangle$ is defined by the noncontextual state. We can write the projector onto this state as

$$\mathcal{I}_{\mathcal{W}} = |\psi_{\text{fixed}}\rangle\langle\psi_{\text{fixed}}| \otimes \mathcal{I}_{(n-|\mathcal{W}^Z|)}, \quad (12)$$

where $\mathcal{I}_{(n-|\mathcal{W}^Z|)}$ is the identity operator acting on the $(n - |\mathcal{W}^Z|)$ qubits not fixed by the single-qubit P_v stabilizers. The action on a general state $|\phi\rangle$ is

$$\mathcal{I}_{\mathcal{W}}|\phi\rangle = |\psi_{\text{fixed}}\rangle\langle\psi_{\text{fixed}}|\phi\rangle \otimes |\phi\rangle_{(n-|\mathcal{W}^Z|)}, \quad (13)$$

where $\mathcal{I}_{\mathcal{W}}$ has only fixed the state of qubits v and thus each stabilizer P_v removes 1 qubit from the problem. As the states of these qubits are fixed, the expectation values of the single-qubit Pauli matrices indexed on qubits v are known. Thus the Pauli operators in the rotated Hamiltonian $H'_{\text{full}} = U_{\mathcal{W}}^\dagger H_{\text{full}} U_{\mathcal{W}}$ acting on these qubits can be updated accordingly, and the Pauli matrices on qubits v can be dropped. Any term in the rotated Hamiltonian that anticommutes with a fixed generator P_v is forced to have an expectation value of zero and can be completely removed from the problem Hamiltonian. The resultant Hamiltonian acts on $|\mathcal{W}^Z|$ fewer qubits. We denote this operation as $H_{\text{full}} \mapsto H_{\text{full}}^{\mathcal{W}} = Q_{\mathcal{W}}^\dagger U_{\mathcal{W}}^\dagger H_{\text{full}} U_{\mathcal{W}} Q_{\mathcal{W}}$. The noncontextual approximation will be stored in the identity term of the problem and therefore does not need to be tracked separately.

The choice of which stabilizer eigenvalues to fix (i.e., what is included in \mathcal{W}) and which to allow to vary remains an open question of the CS-VQE algorithm. The number of possible stabilizer combinations will be $\sum_{i=1}^{|\mathcal{W}_{\text{all}}|} \binom{|\mathcal{W}_{\text{all}}|}{i} = 2^{|\mathcal{W}_{\text{all}}|} - 1$. Rather than searching over all $2^{|\mathcal{W}_{\text{all}}|} - 1$ combinations of stabilizers to fix, in this paper we use the heuristic given in Ref. [14]. This begins at the full noncontextual approximation, where \mathcal{W} contains all possible stabilizers. We then add a qubit to the quantum correction, by removing an operator from \mathcal{W} and greedily choosing each pair that gives the lowest ground-state energy estimate [49]. Alternative strategies on how to do this remain an open question of CS-VQE. A possible way to approach this problem is to look at the priority of different terms in H_{con} [49]. Note that the quality of the approximation is sensitive to which stabilizers are included in \mathcal{W} . When fewer stabilizers are considered (included in \mathcal{W}), the resultant rotated Hamiltonians will act on more qubits and approximate the true ground-state energy better.

In Ref. [14], Kirby *et al.* construct R as a sequence of rotations (exponentiated Pauli operators) defined by $A(\vec{r})$ as in the unitary partitioning method [15,29]. We denote this

operation R_S . The Supplemental Material gives the full definition of this operator [40]. If R_S is considered as just an arbitrary sequence of exponentiated Pauli operator rotations, then the transformation $H_{\text{full}} \mapsto H_{\text{full}}^{R_S} = R_S H_{\text{full}} R_S^\dagger$ results in an operator whose terms have increased by a factor of $O(2^N)$, where N is the number of cliques defined from \mathcal{T} [14]. This presents a possible roadblock for the CS-VQE algorithm, as classically precomputing $U_{\mathcal{W}}^\dagger H_{\text{full}} U_{\mathcal{W}}$ could cause the number of terms to exponentially increase. We give a further analysis of this in the Supplemental Material [40]. Additional structure between R_S and H_{full} can make the base of the exponent slightly lower; however, the scaling still remains exponential in the number of qubits n , where $|\mathcal{A}| \leq 2n + 1$ [42]. The only case in which there is not an exponential increase in terms is for the trivial instance that R_S commutes with H_{full} . In the next section, we provide an alternative construction of R via a linear combination of unitaries (LCU) that results in only a quadratic increase in the number of terms of the Hamiltonian when transformed. This avoids the need to apply the unitary partitioning operator R (via a sequence of rotations) coherently in the quantum circuit after the ansatz circuit, which was proposed in Ref. [14].

F. Linear-combination-of-unitaries construction of R

In the unitary partitioning method [15,29], it was shown that R could also be built as a linear combination of Pauli operators [15,30]. We provide the full construction in the Supplemental Material [40]. We denote the operator as R_{LCU} . Rotating a general Hamiltonian H_{full} by this operation R_{LCU} results in

$$R_{\text{LCU}} H_{\text{full}} R_{\text{LCU}}^\dagger = \sum_i^{|H_{\text{full}}|} (\mu_i) P_i + \sum_j^{|A|-1} \sum_{\substack{i \\ \forall \{P_j, P_k, P_i\} = 0}}^{|H_{\text{full}}|} \mu_{ij} P_j P_k P_i + \sum_j^{|A|-1} \sum_i^{|H_{\text{full}}|} \sum_{\substack{l > j \\ \forall \{P_i, P_j, P_l\} = 0}}^{|A|-1} \mu_{ijl} P_i P_j P_l. \quad (14)$$

The Pauli operators P_j , P_k , and P_l are operators in \mathcal{A} ; further details are covered in the Supplemental Material [40]. Overall, this unitary transformation causes the number of terms in the Hamiltonian to scale as $O(|H_{\text{full}}| \cdot |\mathcal{A}|^2)$. This scaling is quadratic in the size of \mathcal{A} , and as $|\mathcal{A}| \leq 2n + 1$ [42], the number of terms in the rotated system will at worst scale quadratically with the number of qubits n . In a different context, this scaling result was also obtained for involutory linear combinations of entanglers [50]. Overall, unlike the sequence-of-rotations approach, this non-Clifford operation does not cause the number of terms in a Hamiltonian to increase exponentially.

The transformation given in Eq. (14), $H_{\text{full}} \mapsto H_{\text{full}}^{\text{LCU}} = R_{\text{LCU}}^\dagger H_{\text{full}} R_{\text{LCU}}$, is performed classically in CS-VQE. This is efficient to do because it just involves Pauli operator multiplication, which can be done symbolically or via a symplectic approach [51]. This operation could be applied within the

quantum circuit. However, in contrast to the deterministic sequence-of-rotations approach, this implementation would be probabilistic as it requires postselection on an ancillary register [15,30,52,53]. Amplitude amplification techniques could improve this but would require further coherent resources [54–57]. Performing this transformation in a classical preprocessing step therefore reduces the coherent resources required and at worst increases the number of terms needing measuring quadratically with respect to the number of qubits.

G. CS-VQE implementation

In Ref. [14], $U_{\mathcal{W}}(\vec{q}, \vec{r})$ was fixed to include all the stabilizers of the noncontextual ground state $\mathcal{W} \equiv \mathcal{W}_{\text{all}}$ [Eq. (8)], rather than possible subsets $\mathcal{W} \subseteq \mathcal{W}_{\text{all}}$. The whole Hamiltonian was mapped according to $H_{\text{full}} \mapsto H'_{\text{full}} = U_{\mathcal{W}_{\text{all}}}^\dagger H_{\text{full}} U_{\mathcal{W}_{\text{all}}}$. In general, $A(\vec{r}) \in \mathcal{W}$, and $U_{\mathcal{W}_{\text{all}}}$ will therefore normally include the unitary partitioning operator R . The problem with this approach is that the unitary R is not a Clifford operation and the transformation can cause the number of terms in the Hamiltonian to increase. This increase is exponential if R_S is used and quadratic if R_{LCU} is employed. As this step can generate more terms, R should only be included in $U_{\mathcal{W}}$ if the eigenvalue of $A(\vec{r})$ is fixed to +1; otherwise it is a redundant operation as the spectrum of the operator rotated by R is unchanged. We therefore modify the CS-VQE algorithm to construct $U_{\mathcal{W}}$ from the CS-VQE noncontextual generator eigenvalues that are fixed. This means that $\mathcal{W} \subseteq \mathcal{W}_{\text{all}}$ and ensures that the number of terms can only increase if the eigenvalue of $A(\vec{r})$ is fixed.

III. NUMERICAL RESULTS

We describe the method in Sec. III A and then split our results into the two remaining sections. First, we explore a toy problem, showing the steps of the CS-VQE algorithm. We show how classically applying R without fixing the eigenvalue of $A(\vec{r})$ to +1 can unnecessarily increase the number of terms in a Hamiltonian without changing its spectrum. Finally, in Sec. III C we apply measurement reduction combined with CS-VQE to a set of electronic structure Hamiltonians and show that this can significantly reduce the number of terms requiring separate measurement. The raw data for these results are supplied in the Supplemental Material [40].

A. Method

We investigated the same electronic structure Hamiltonians considered in the original CS-VQE paper [14]. All molecules considered had a multiplicity of 1 and thus a singlet ground state. The same qubit tapering was performed to remove the \mathbb{Z}_2 symmetries [58]. For each tapered Hamiltonian, we generate a set of reduced Hamiltonians $\{Q_{\mathcal{W}}^\dagger U_{\mathcal{W}}^\dagger H_{\text{full}} U_{\mathcal{W}} Q_{\mathcal{W}}\}$, where the size of \mathcal{W} varies from 1 to $|\mathcal{W}_{\text{all}}|$, representing differing noncontextual approximations, as summarized in Sec. II G. To generate the different CS-VQE Hamiltonians, we modify the original CS-VQE source code used in Refs. [14,59]. The code was modified to implement the unitary partitioning step of CS-VQE if and only if the eigenvalue of $A(\vec{r})$ was fixed. This ensured that the number of terms

in the rotated Hamiltonian did not increase unnecessarily, as described in Sec. II G.

For each electronic structure Hamiltonian generated in this way, we then apply the unitary partitioning measurement reduction scheme to further reduce the number of terms requiring separate measurement [15,29,30]. Partitioning into anticommuting sets was performed using NETWORKX [60]. A graph of the qubit Hamiltonian is built, where nodes represent Pauli operators and edges are between nodes that commute. A graph coloring can be used to find the anticommuting cliques of the graph. This searches for the minimum number of colors required to color the graph, where no neighbors of a node can have the same color as the node itself. The largest-first coloring strategy in NETWORKX was used in all cases [60,61].

We calculate the ground-state energy of each Hamiltonian in this paper by directly evaluating the lowest eigenvalues. This was achieved by diagonalizing them on a conventional computer.

B. Toy example

We consider the qubit Hamiltonian

$$\begin{aligned} H = & 0.6 IIYI + 0.7 XYXI + 0.7 XZXI + 0.6 XZZI \\ & + 0.1 YXYI + 0.7 ZZZI + 0.5 IIIZ + 0.1 XXXI \\ & + 0.5 XXYI + 0.2 XXZI + 0.2 YXXI + 0.2 YYZI \\ & + 0.1 YZXI + 0.1 ZYYI \end{aligned} \quad (15)$$

and use it to exemplify the steps of the CS-VQE algorithm. The results are reported to three decimal places, and full numerical details can be found in the Supplemental Material [40].

Following the CS-VQE procedure [14], we first split the Hamiltonian into its contextual and noncontextual parts [Eq. (2)]:

$$\begin{aligned} H_{\text{noncon}} = & 0.5 \underbrace{III}_{Z} \\ & + 0.7 XZXI + 0.7 ZZZI \\ & + 0.1 YXYI + 0.6 IIYI \\ & + 0.7 \underbrace{XYXI + 0.6 XZZI}_{T}, \end{aligned} \quad (16a)$$

$$\begin{aligned} H_{\text{con}} = & 0.1 XXXI + 0.5 XXYI + 0.2 XXZI \\ & + 0.2 YXXI + 0.2 YYZI + 0.1 YZXI \\ & + 0.1 ZYYI. \end{aligned} \quad (16b)$$

Each row after the first in Eq. (16a) is a clique of \mathcal{T} . From here, we define the set \mathcal{R} [Eq. (4)]:

$$\mathcal{R} = \underbrace{\{YIYI, IXYI, IIIZ\}}_{\mathcal{G}} \cup \underbrace{\{XZXI, YXYI, XYXI\}}_{\{P_0^{(j)} | j=0,1,\dots,N-1\}}. \quad (17)$$

Note how different combinations of the operators in Eq. (17) allow all the operators in H_{noncon} [Eq. (16a)] to be inferred under the Jordan product, defined as $P_a \circ P_b = \frac{1}{2}(P_a P_b + P_b P_a)$. Basically, the Jordan product is equal to the regular matrix product if the operators commute and equal to zero if the operators anticommute. Next the noncontextual problem was solved.

The expectation value for H_{noncon} can be induced, by setting the expectation values of operators in \mathcal{R} [Eq. (17)], as the Pauli operators in H_{noncon} are generated by \mathcal{R} under the Jordan product. The expectation value of each operator in H_{noncon} can therefore be inferred without contradiction. To find the ground state of H_{noncon} , we checked all possible ± 1 expectation values for each G_j ($2^3 = 8$ possibilities). For each possible ± 1 combination, the energy was minimized with respect to the unit vector \vec{r} , which sets the expectation value for each $\langle P_0^{(j)} \rangle = r_j$. The vector (\vec{q}, \vec{r}) that was found to give the lowest energy defines the noncontextual ground state. In this case the ground state is

$$(\underbrace{-1, +1, -1}_{\vec{q}_0}, \underbrace{+0.253, -0.658, -0.709}_{\vec{r}_0}). \quad (18)$$

This noncontextual state defines the operator $\mathcal{A}(\vec{r}_0)$:

$$\mathcal{A}(\vec{r}_0) = 0.253 YXYI - 0.658 XYXI - 0.709 XZXI. \quad (19)$$

From this we can write $\mathcal{R}_{\vec{r}}$ [Eq. (7)]:

$$\mathcal{R}_{\vec{r}} \equiv \{\mathcal{A}(\vec{r}_0)\} \cup \{YIYI, IXYI, IIIZ\}. \quad (20)$$

To map $\mathcal{A}(\vec{r}_0)$ to a single Pauli operator, we use unitary partitioning [15,29,30]. The required unitary can be constructed as either a sequence of rotations [15],

$$R_S = e^{-1i \cdot 0.788 \cdot ZYZI} \cdot e^{+1i \cdot 1.204 \cdot ZZZI}, \quad (21)$$

or linear combination of unitaries [15],

$$R_{\text{LCU}} = 0.792 IIII + 0.416i ZZZI - 0.448i ZYZI. \quad (22)$$

These operators perform the following reduction: $R_S \mathcal{A}(\vec{r}_0) R_S^\dagger = R_{\text{LCU}} \mathcal{A}(\vec{r}_0) R_{\text{LCU}}^\dagger = YXYI$.

If the eigenvalue of $\mathcal{A}(\vec{r}_0)$ is fixed, then we should consider $\mathcal{R}_{\vec{r}}$ [Eq. (20)] under the unitary transform R_{LCU} or R_S [Eq. (7)]:

$$\begin{aligned} \mathcal{R}' &= R_{S/\text{LCU}} \mathcal{A}(\vec{r}) R_{S/\text{LCU}}^\dagger \cup \{YIYI, IXYI, IIIZ\} \\ &= \{YXYI\} \cup \{YIYI, IXYI, IIIZ\}. \end{aligned} \quad (23)$$

Equations (18), (20), and (23) define the noncontextual stabilizers:

$$\begin{aligned} \mathcal{W}_{\text{all}} &\equiv \{+1 \mathcal{A}(\vec{r}_0), -1 YIYI, +1 IXYI, -1 IIIZ\}, \\ \mathcal{W}'_{\text{all}} &\equiv \{+1 YXYI, -1 YIYI, +1 IXYI, -1 IIIZ\}. \end{aligned} \quad (24)$$

Next, we define different $U_{\mathcal{W}}$ [Eq. (10)], depending on which stabilizers \mathcal{W} we wish to fix. For this problem we found the optimal ordering of which stabilizers to fix to be $\{-1 YIYI, +1 IXYI, +1 \mathcal{A}(\vec{r}_0), -1 IIIZ\}$ followed by $\{+1 IXYI, -1 IIIZ, +1 \mathcal{A}(\vec{r}_0)\}$ followed by $\{+1 IXYI, -1 IIIZ\}$ followed by $\{-1 IIIZ\}$. This was achieved by a brute-force search over all $\sum_{i=1}^{|\mathcal{W}_{\text{all}}|} \binom{|\mathcal{W}_{\text{all}}|}{i} = 2^4 - 1 = 15$ possibilities for \mathcal{W} .

The members of the resulting set of four different \mathcal{W} each represent different noncontextual approximations. These give four different $U_{\mathcal{W}}$ built according to Eq. (10). The full definition of each operator is given in the Supplemental Material [40].

Taking a specific example, for $\mathcal{W} = \{+1 IXYI, -1 IIIZ, +1 \mathcal{A}(\vec{r}_0)\}$ we define $U_{\mathcal{W}}^\dagger$ [Eq. (10)]. This operator transforms \mathcal{W}

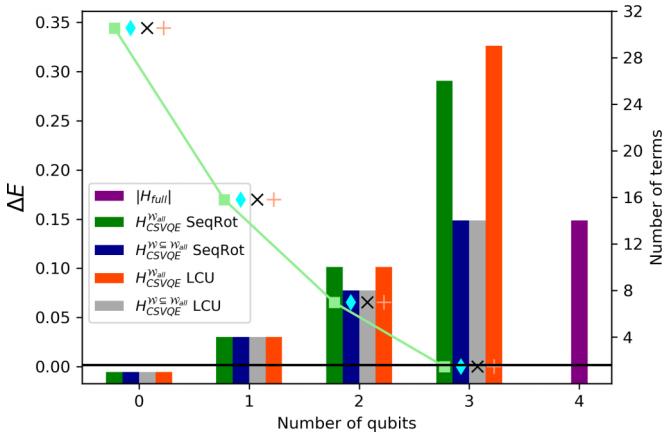


FIG. 1. Ground-state energy and the number of terms of different contextual subspace projected Hamiltonians generated in CS-VQE. Each Hamiltonian has been transformed as $Q_{\mathcal{W}}^{\dagger} U_{\mathcal{W}}^{\dagger} H U_{\mathcal{W}} Q_{\mathcal{W}}$, apart from the 4-qubit case, which is the full H . The scatterplot is associated with the left-hand y axis and gives the energy error as $\Delta E = |E_{\text{approx}} - E_{\text{true}}|$. The bar chart gives the number of terms in each Hamiltonian and is associated with the right-hand y axis. From left to right the following generators are fixed: $\{YIYI, IXYI, IIIZ, \mathcal{A}(\vec{r}_0)\}$, $\{IXYI, IIIZ, \mathcal{A}(\vec{r}_0)\}$, $\{IXYI, IIIZ\}$, $\{IIIZ\}$, and $\{\}$. The $\mathcal{W} = \{\}$ case represents standard full VQE over the full problem. The 0-qubit case presents the scenario where the problem is fully noncontextual and no quantum correction is made. The full details as to how each Hamiltonian is built are provided in the Supplemental Material [40]. The horizontal black line indicates an absolute error of 1.6×10^{-3} . SeqRot, sequence of rotations.

as $\mathcal{W}^Z = U_{\mathcal{W}}^{\dagger} \mathcal{W} U_{\mathcal{W}} = \{+IZII, -IIIZ, +IIZI\}$. The eigenvalues of the operators in \mathcal{W}^Z are fixed by the noncontextual state to be $\langle IZII \rangle = +1$, $\langle IIZI \rangle = +1$, $\langle IIIZ \rangle = -1$. This defines the projector:

$$\begin{aligned} Q_{\mathcal{W}} &= (\underbrace{|0\rangle\langle 0| + |1\rangle\langle 1|}_{\mathcal{I}_{(n-|\mathcal{W}^Z|)}} \otimes \underbrace{|0\rangle\langle 0| \otimes |0\rangle\langle 0| \otimes |1\rangle\langle 1|}_{|\psi_{\text{fixed}}\rangle} \\ &= I \otimes |001\rangle\langle 001|. \end{aligned} \quad (25)$$

The reduced Hamiltonian is therefore

$$\begin{aligned} H \mapsto H_{\mathcal{W}}^{\text{LCU}} &= Q_{\mathcal{W}}^{\dagger} U_{\mathcal{W}}^{\dagger} (H_{\text{full}} U_{\mathcal{W}}^{\text{LCU}}) Q_{\mathcal{W}} \\ &= -1.827 I - 0.414 X - 0.292 Z + 0.648 Y. \end{aligned} \quad (26)$$

The Supplemental Material gives further details about this operation and provides the specifics for the other projected Hamiltonians [40].

Overall, four Hamiltonians are generated, representing different levels of approximation, that act on 0, 1, 2, and 3 qubits, respectively. The 4-qubit case represents the standard VQE on the full Hamiltonian. Figure 1 summarizes the error ΔE of each of these compared with the true ground-state energy (scatterplot). The number of terms in each Hamiltonian is given by the bar chart. The green and orange results have $\mathcal{W} \equiv \mathcal{W}_{\text{all}}$ for all cases and represent the old CS-VQE implementation. For these results, the 3- and 4-qubit Hamiltonians have an increased number of terms due to $R_{S/\text{LCU}}$ being implemented, even though the eigenvalue of $\mathcal{A}(\vec{r}_0)$ is not being

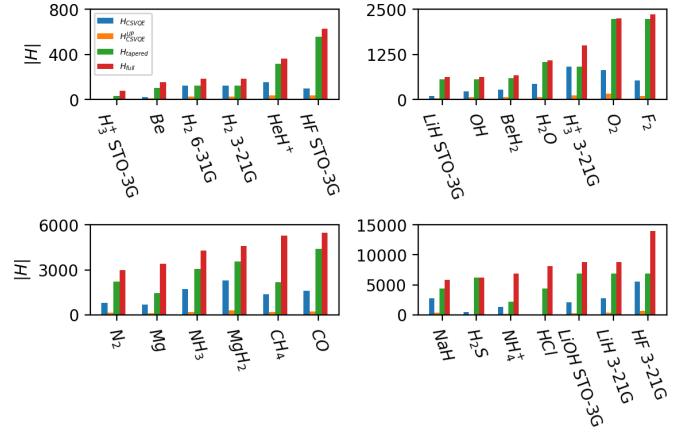


FIG. 2. Number of Pauli operators requiring separate measurement to determine the ground-state energy of a particular molecular Hamiltonian to within chemical accuracy. For each molecule the full Hamiltonian, tapered Hamiltonian, CS-VQE, and CS-VQE with unitary partitioning (UP) measurement reduction applied are given. Full numerical numerical details of each are provided in the Supplemental Material [40]. The size of the Hamiltonian for LiH (3-21G singlet) with measurement reduction applied is different for the sequence-of-rotations and LCU unitary partitioning methods. This is an artifact of the graph color heuristic finding different anticommuting cliques in the CS-VQE Hamiltonian.

fixed to $+1$. On the other hand, the gray and blue results in Fig. 1 build $U_{\mathcal{W}}$ according to Eq. (10), where $\mathcal{W} \subseteq \mathcal{W}_{\text{all}}$. This approach ensures that $R_{S/\text{LCU}}$ is only applied when necessary.

C. Measurement reduction

Figures 2 and 3 summarize the results of applying the unitary partitioning measurement reduction strategy to a set of electronic structure Hamiltonians. We report the number of terms and number of qubits in each Hamiltonian required to achieve chemical accuracy compared with the origi-

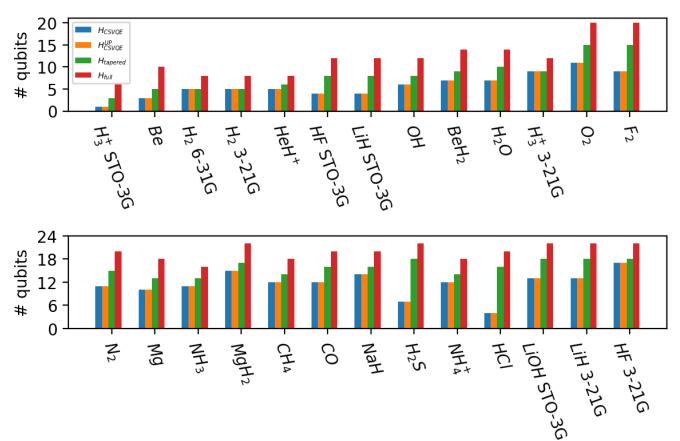


FIG. 3. Number of qubits required to simulate different electronic structure Hamiltonians in order to achieve chemical accuracy. For each molecule the full Hamiltonian, tapered Hamiltonian, CS-VQE, and CS-VQE with unitary partitioning measurement reduction applied are given. Numerical details for each result are provided in the Supplemental Material.

nal problem. The Supplemental Material [40] gives further information about each result, where the different levels of noncontextual approximation are shown. As previously discussed in Ref. [14], even though CS-VQE in general is an approximate method, chemical accuracy can still be achieved using significantly fewer qubits. Applying unitary partitioning on top of the reduced CS-VQE Hamiltonians required to achieve chemical accuracy can further reduce the number of terms by roughly an order of magnitude. This is consistent with the previous results in Ref. [30].

To actually obtain a measurement reduction, one needs to show that the number of measurements required to measure the energy of a molecular system, to a certain precision ϵ , is reduced. Currently, Fig. 2 only shows that we have reduced the number of Pauli terms being measured. We have not commented on the variance. In the Supplemental Material [40], we prove that simultaneous measurement of normalized anticommuting cliques can never do worse than performing no measurement reduction and will more often than not give an improvement. The proof given is state independent. There are other measurement strategies based on grouping techniques, such as splitting a Hamiltonian into commuting or qubit-wise commuting cliques [16,17,21,24,25]. The measurement reduction obtained from these methods is more complicated, as the covariance of operators within a clique must be carefully accounted for [24,62]. This is one of the reasons we do not analyze the performance of these strategies in this paper. Many other measurement methods have also been proposed [18–20,22,23,25–28,63,64], and their effect on the number of measurements would be interesting to investigate.

In Table I, we report the upper bound on the gate count required to implement measurement reduction as a sequence of rotations. The LCU method would require ancilla qubits, and analysis of the circuit depth is more complicated. Further analysis can be found in Ref. [30]. The number of extra coherent resources required to implement unitary partitioning measurement reduction is proportional to the size of each anticommuting clique a Hamiltonian is split into [15,30]. The sequence-of-rotations circuit depth scales as $O(N_s(|C| - 1))$ single-qubit and $O(N_s(|C| - 1))$ CNOT gates, where N_s is the number of system qubits and $|C|$ is the size of the anticommuting clique being measured. Table I reports the gate count upper bound for the largest anticommuting clique of a given CS-VQE Hamiltonian. We do not consider possible circuit simplifications, such as gate cancellations. To decrease the depth of the quantum circuit required for practical application, we suggest finding nonoptimal clique covers; for example, if anticommuting cliques are fixed to a size of 2, the resources required to perform R_S are experimentally realistic for current and near-term devices, as only $O(N_s)$ single-qubit and $O(N_s)$ CNOT gates are required [30].

The heuristic used to determine the operators in H_{noncon} selected terms in the full Hamiltonian greedily by coefficient magnitude, while keeping the set noncontextual [42]. The Hamiltonians studied here had weights dominated by diagonal Pauli operators, as the Hartree-Fock approximation accounts for most of the energy. This heavily constrains the operators allowed in \mathcal{A} . For the electronic structure Hamiltonians considered in this paper, we found in all cases that $|\mathcal{A}| = 2$. In general, we do expect more commuting terms in H_{noncon}

TABLE I. Gate requirements to implement R as a sequence of rotations in the unitary partitioning measurement reduction step. The square tuple gives the upper bound on the number of single-qubit and CNOT gates required: [single, CNOT]. These resource requirements are based on the largest anticommuting clique of each Hamiltonian, as these have the largest circuit requirements for R_S .

Molecule	Basis	Number of gates for R_S
BeH ₂	STO-3G	[90, 72]
Mg	STO-3G	[189, 162]
H ₃ ⁺	3-21G	[209, 176]
O ₂	STO-3G	[184, 160]
OH ⁻	STO-3G	[104, 80]
CH ₄	STO-3G	[325, 286]
Be	STO-3G	[14, 8]
NH ₃	STO-3G	[299, 260]
H ₂ S	STO-3G	[120, 96]
H ₂	3-21G	[66, 48]
HF	3-21G	[735, 672]
F ₂	STO-3G	[133, 112]
HCl	STO-3G	[36, 24]
HeH ⁺	3-21G	[88, 64]
MgH ₂	STO-3G	[403, 364]
CO	STO-3G	[325, 286]
LiH	STO-3G	[36, 24]
N ₂	STO-3G	[207, 180]
NaH	STO-3G	[493, 442]
H ₂ O	STO-3G	[120, 96]
H ₃ ⁺	STO-3G	[3, 0]
LiOH	STO-3G	[378, 336]
LiH	3-21G	[459, 408]
H ₂	6-31G	[66, 48]
NH ₄ ⁺	STO-3G	[325, 286]
HF	STO-3G	[36, 24]

than anticommuting terms. This is because there are more possible commuting Pauli operators defined on n qubits compared with anticommuting operators (2^n vs $2n + 1$). \mathcal{G} will therefore in general be the larger contributor to the superset \mathcal{R} [Eq. (4)].

In Fig. 2, the CS-VQE bars have not been split into two for the case when R is constructed as R_{LCU} or R_S . This is due to $|\mathcal{A}|$ being 2 in all cases, which is the special case when these operators (R_{LCU} and R_S) end up being identical. In this instance, R has the form $R = \alpha I + i\beta P$, and thus the number of terms will only increase for every term in the Hamiltonian that P anticommutes with. However, in general, $|\mathcal{A}|$ will be greater than 2, and the effect of R can dramatically affect the number of terms in the resultant rotated Hamiltonian. We observe this in Fig. 1 of the toy example, where the 2- and 3-qubit CS-VQE Hamiltonians have had $U_{\mathcal{W}_{\text{all}}}$ applied to them even though the eigenvalue of $A(\vec{r})$ is not fixed. In that example, for the 3-qubit approximation the sequence-of-rotations rotated Hamiltonian (green) actually has fewer terms than the LCU rotated operator (orange). This result is an artifact of the small problem size. In the Supplemental Material we show that the scaling will favor the LCU implementation, where the number of terms in a Hamiltonian can only increase quadratically,

not exponentially, when performing the unitary partitioning rotation as a LCU rather than a sequence of rotations [40].

In the Supplemental Material, we show the convergence of CS-VQE at different noncontextual approximations. The results illustrate that CS-VQE can converge to below chemical accuracy well before the case when no noncontextual approximation is made (full VQE). Results beyond convergence are included to show the different possible levels of approximation. In practice, knowledge of the true ground-state energy is not known *a priori*, and so using chemical precision to motivate the noncontextual approximation will not be possible. In this setting, a way to approach quantum advantage is to note that CS-VQE is a variational method. The quantum resources required can be expanded until the energy obtained by CS-VQE is lower than that coming from the best possible classical method. At this point, either the algorithm can be terminated or further contextual corrections can be added until the energy converges, at which point the algorithm should be stopped.

IV. CONCLUSION

The work presented here shows that combining the unitary partitioning measurement reduction strategy with the CS-VQE algorithm can further reduce the number of terms in the projected Hamiltonian requiring separate measurement by roughly an order of magnitude for a given molecular Hamiltonian. The number of qubits needed to achieve chemical accuracy in most cases was also dramatically decreased, for example, the H₂S (STO-3G singlet) problem was reduced to 7 qubits from 22.

We also improve two parts of the CS-VQE algorithm. First, we avoid having to apply the unitary partitioning operator R after the ansatz, which prevents the potential exponential increase in the number of Pauli operators of the CS-VQE Hamiltonian caused by classically computing the non-Clifford rotation of the full Hamiltonian when R is defined as a sequence of rotations [14,15]. We show that applying this operation as a linear combination of unitaries [15], $H_{\text{full}} \mapsto H_{\text{full}}^{\text{LCU}} = R_{\text{LCU}}^\dagger H_{\text{full}} R_{\text{LCU}}$, results in the number of terms at worst increasing quadratically with the number of qubits. This result makes classically precomputing this transformation tractable, and R no longer needs to be performed coherently after the ansatz. Secondly, we define the unitary $U_{\mathcal{W}}$, which maps each stabilizer in \mathcal{W}_{all} [Eq. (9)] to a distinct single-qubit Pauli matrix, according to which stabilizer eigenvalues are fixed by the noncontextual state. This ensures that the non-Clifford rotation required by CS-VQE is only applied when necessary and also reduces the number of redundant Clifford operations that are classically performed.

There are still several open questions for the CS-VQE algorithm. We summarize a few here. (1) What is the best optimization strategy to use when minimizing the energy over (\vec{q}, \vec{r}) in the classical noncontextual problem? (2) What heuristic is best to construct the largest $|H_{\text{noncon}}|$? (3) How can we efficiently determine which noncontextual stabilizers to fix while maintaining low errors? In this paper, the size of each electronic structure problem allowed us to classically compute the ground-state energies at each step, but if this is not possible, then VQE calculations would be required.

However, as each run requires fewer qubits and decreases the number of terms requiring separate measurement, this approach may overall still be less costly than performing VQE over the whole problem, especially when combined with further measurement reduction strategies. (4) What are the most important terms to include in H_{con} or, equivalently, in H_{noncon} ? Currently, it is not known whether $|H_{\text{noncon}}|$ should be maximized or whether selecting high-priority terms [49] from the whole Hamiltonian results in a better approximation for a given problem. We leave these questions to future work.

We have written an open-source CS-VQE code that includes all the updated methodology discussed in this paper. We welcome readers to make use of this code, which is freely available on GitHub [65].

ACKNOWLEDGMENTS

A.R. and T.W. acknowledge support from the Unitary Fund and the Engineering and Physical Sciences Research Council (Grants No. EP/L015242/1 and No. EP/S021582/1, respectively). T.W. also acknowledges support from CBKSci-Con Ltd., Atos, Intel, and Zapata. W.M.K. and P.J.L. acknowledge support by the NSF STAQ project (Grant No. PHY-1818914). W.M.K. acknowledges support from the National Science Foundation, Grant No. DGE-1842474. P.V.C. is grateful for funding from the European Commission for VECMA (800925) and EPSRC for SEAVER (Grant No. EP/W007711/1).

APPENDIX: PERES-MERMIN SQUARE

The “Peres-Mermin square” (PM square) [37,39] involves the construction of nine measurements arranged in a square. In this Appendix we follow the construction given in Ref. [33]. Each measurement has only two possible outcomes (dichotomic), +1 and -1. In a realistic interpretation, performing each measurement on an object reveals whether the property is present (+1) or absent (-1), yielding nine properties.

We take three measurements along a column or row to form a “context”: a set of measurements whose values can be jointly measured, i.e., the observables commute and thus share a common eigenbasis. Table II gives an example.

In a classical (noncontextual) model for this system, the nine measurements $\{IZ, ZI, ZZ, XI, IX, XZ, ZX, YY\}$ can be assigned a definite value independent of the context the measurement is obtained in. For example, if all measurements are assigned +1 in Table II, then $c_0 = c_1 = c_2 = r_0 = r_1 = r_2 = +1$, and six positive products are obtained. If a single entry in Table II is changed, it will affect two products (a row

TABLE II. Example Peres-Mermin square of nine possible observables for a physical system, where each measurement can be assigned a ± 1 value.

	c_0	c_1	c_2
r_0	IZ	ZI	ZZ
r_1	XI	IX	XX
r_2	XZ	ZX	YY

and column product). We consider the following equation in this setting:

$$\begin{aligned} \langle PM \rangle &\equiv \langle IZ \cdot ZI \cdot ZZ \rangle + \langle XI \cdot IX \cdot XX \rangle \\ &+ \langle XZ \cdot ZX \cdot YY \rangle + \langle IZ \cdot XI \cdot XZ \rangle \\ &+ \langle ZI \cdot IX \cdot ZX \rangle - \langle ZZ \cdot XX \cdot YY \rangle \\ &= r_0 + r_1 + r_2 + c_0 + c_1 - c_2. \end{aligned} \quad (A1)$$

We find that classically, we get an inequality, $\langle PM \rangle \leq 4$. We reiterate that this is the setting of eight $+1$ assignments and a single -1 assignment. This inequality is saturated when the -1 value is assigned to one of the observables in the last column of Table II.

The significance of this inequality is that it can be violated by quantum systems. Thinking of this in a quantum setting, the operators in the rows and columns of Table II commute. If we multiply along the rows and columns, we get $+II$, apart from the last column, where $c_2 = -II$ (see Table III). This is the case regardless of what quantum state is considered. Using the expectation values of the product of these operators in Eq. (A1), we find $\langle PM \rangle = 6$, violating the classical bound.

Classically, Eq. (A1) is bounded as $\langle PM \rangle \leq 4$ due to the assumption that the nine observables of the object can be assigned a value consistently. Violation of this bound implies that either the value assignment must depend on which context

TABLE III. Example Peres-Mermin square of nine Hermitian operators, all with ± 1 eigenvalues, representing observables.

	$\langle +II \rangle = +1$	$\langle +II \rangle = +1$	$\langle -II \rangle = -1$
$\langle +II \rangle = +1$	IZ	ZI	ZZ
$\langle +II \rangle = +1$	XI	IX	XX
$\langle +II \rangle = +1$	XZ	ZX	YY

(row or column) the observable appears in or there is no value assignment. This phenomenon is known as quantum contextuality [33].

In VQE, a Hamiltonian is defined by a linear combination of Pauli operators. The expectation value is obtained by measuring each Pauli operator in a separate experiment and combining the results. Different groups of commuting operators form contexts. In general, there will be incompatible contexts where it is impossible to consistently assign joint outcomes. In other words, different inference relations will lead to contradictions. Outcomes assigned to individual measurements are therefore context dependent, and the problem is contextual. If not, then the problem is noncontextual, and a noncontextual (classical) hidden-variable model can be used to solve such systems.

[1] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Macmillan, New York, 2012).

[2] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).

[3] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity, *Phys. Rev. X* **8**, 041015 (2018).

[4] V. E. Elfving, B. W. Broer, M. Webber, J. Gavartin, M. D. Halls, K. P. Lorton, and A. Bochevarov, How will quantum computers provide an industrially relevant computational advantage in quantum chemistry? [arXiv:2009.12472](https://arxiv.org/abs/2009.12472).

[5] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble, and R. C. Pooser, Quantum chemistry as a benchmark for near-term quantum computers, *npj Quantum Inf.* **5**, 99 (2019).

[6] A. Y. Kitaev, Quantum measurements and the Abelian stabilizer problem, [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).

[7] 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 *et al.*, Scalable Quantum Simulation of Molecular Energies, *Phys. Rev. X* **6**, 031007 (2016).

[8] H. Mohammadbagherpoor, Y.-H. Oh, A. Singh, X. Yu, and A. J. Rindos, Experimental challenges of implementing quantum phase estimation algorithms on IBM quantum computer, [arXiv:1903.07605](https://arxiv.org/abs/1903.07605).

[9] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).

[10] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).

[11] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, Variational quantum linear solver, [arXiv:1909.05820](https://arxiv.org/abs/1909.05820).

[12] A. Eddins, M. Motta, T. P. Gujarati, S. Bravyi, A. Mezzacapo, C. Hadfield, and S. Sheldon, Doubling the size of quantum simulators by entanglement forging, *PRX Quantum* **3**, 010309 (2022).

[13] W. J. Huggins, B. A. O’Gorman, N. C. Rubin, D. R. Reichman, R. Babbush, and J. Lee, Unbiasing fermionic quantum Monte Carlo with a quantum computer, *Nature (London)* **603**, 416 (2022).

[14] W. M. Kirby, A. Tranter, and P. J. Love, Contextual subspace variational quantum eigensolver, *Quantum* **5**, 456 (2021).

[15] A. Zhao, A. Tranter, W. M. Kirby, S. F. Ung, A. Miyake, and P. J. Love, Measurement reduction in variational quantum algorithms, *Phys. Rev. A* **101**, 062322 (2020).

[16] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).

[17] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover, *J. Chem. Phys.* **152**, 124114 (2020).

[18] A. F. Izmaylov, T.-C. Yen, and I. G. Ryabinkin, Revising the measurement process in the variational quantum eigensolver:

is it possible to reduce the number of separately measured operators? *Chem. Sci.* **10**, 3746 (2019).

[19] J. Cotler and F. Wilczek, Quantum Overlapping Tomography, *Phys. Rev. Lett.* **124**, 100401 (2020).

[20] X. Bonet-Monroig, R. Babbush, and T. E. O'Brien, Nearly Optimal Measurement Scheduling for Partial Tomography of Quantum States, *Phys. Rev. X* **10**, 031064 (2020).

[21] P. Gokhale and F. T. Chong, $O(N^3)$ measurement cost for variational quantum eigensolver on molecular Hamiltonians, in *IEEE Transactions on Quantum Engineering*, Vol. 1 (IEEE, 2020), pp. 1–24.

[22] A. Jena, S. Genin, and M. Mosca, Pauli partitioning with respect to gate sets, [arXiv:1907.07859](https://arxiv.org/abs/1907.07859).

[23] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *npj Quantum Inf.* **7**, 23 (2021).

[24] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, M. Suchara, M. Martonosi, and F. T. Chong, Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families, [arXiv:1907.13623](https://arxiv.org/abs/1907.13623).

[25] O. Crawford, B. van Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of Pauli operators in the presence of finite sampling error, *Quantum* **5**, 385 (2021).

[26] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).

[27] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of quantum Hamiltonians with locally-biased classical shadows, *Commun. Math. Phys.* **391**, 951 (2022).

[28] H.-Y. Huang, R. Kueng, and J. Preskill, Efficient Estimation of Pauli Observables by Derandomization, *Phys. Rev. Lett.* **127**, 030503 (2021).

[29] A. F. Izmaylov, T.-C. Yen, R. A. Lang, and V. Verteletskyi, Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method, *J. Chem. Theory Comput.* **16**, 190 (2020).

[30] A. Ralli, P. J. Love, A. Tranter, and P. V. Coveney, Implementation of measurement reduction for the variational quantum eigensolver, *Phys. Rev. Res.* **3**, 033195 (2021).

[31] S. Kochen and E. P. Specker, The problem of hidden variables in quantum mechanics, in *The Logico-algebraic Approach to Quantum Mechanics* (Springer, New York, 1975), pp. 293–328.

[32] C. Budroni, Contextuality, memory cost and non-classicality for sequential measurements, *Philos. Trans. R. Soc. A* **377**, 20190141 (2019).

[33] C. Budroni, A. Cabello, O. Gühne, M. Kleinmann, and J.-Å. Larsson, Kochen-Specker contextuality, *Rev. Mod. Phys.* **94**, 045007 (2022).

[34] N. de Silva, Graph-theoretic strengths of contextuality, *Phys. Rev. A* **95**, 032108 (2017).

[35] A. Cabello, Converting Contextuality into Nonlocality, *Phys. Rev. Lett.* **127**, 070401 (2021).

[36] M. Howard, J. Wallman, V. Veitch, and J. Emerson, Contextuality supplies the ‘magic’ for quantum computation, *Nature (London)* **510**, 351 (2014).

[37] N. D. Mermin, Simple Unified Form for the Major No-Hidden-Variables Theorems, *Phys. Rev. Lett.* **65**, 3373 (1990).

[38] J. S. Bell, On the Einstein Podolsky Rosen paradox, *Phys. Phys. Fiz.* **1**, 195 (1964).

[39] A. Peres, Incompatible results of quantum measurements, *Phys. Lett. A* **151**, 107 (1990).

[40] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.5.013095>. for further information on the CS-VQE algorithm, the scaling analysis of the unitary partitioning rotation used in CS-VQE, the analysis of the variance of the ground-state energy obtained when using unitary partitioning for measurement reduction, and the numerical details of the physical problems presented in this paper.

[41] R. W. Spekkens, Quasi-quantization: Classical statistical theories with an epistemic restriction, in *Quantum Theory: Informational Foundations and Foils*, edited by G. Chiribella and R. W. Spekkens (Springer, Dordrecht, 2016), pp. 83–135.

[42] W. M. Kirby and P. J. Love, Classical simulation of noncontextual Pauli Hamiltonians, *Phys. Rev. A* **102**, 032418 (2020).

[43] R. W. Spekkens, Evidence for the epistemic view of quantum states: A toy theory, *Phys. Rev. A* **75**, 032110 (2007).

[44] R. W. Spekkens, Quasi-quantization: classical statistical theories with an epistemic restriction, in *Quantum Theory: Informational Foundations and Foils*, edited by G. Chiribella and R. W. Spekkens (Springer, Dordrecht, 2016), pp. 83–135.

[45] T. J. Weaving, A. Ralli, W. M. Kirby, A. Tranter, P. J. Love, and P. V. Coveney, A stabilizer framework for contextual subspace variational quantum eigensolver and the noncontextual projection ansatz, *J. Chem. Theory Comput.* (2023), doi: [10.1021/acs.jctc.2c00910](https://doi.org/10.1021/acs.jctc.2c00910).

[46] R. Raussendorf, J. Bermejo-Vega, E. Tyhurst, C. Okay, and M. Zurel, Phase-space-simulation method for quantum computation with magic states on qubits, *Phys. Rev. A* **101**, 012350 (2020).

[47] Note that a stabilizer for a state leaves it unchanged. For example, if O stabilizes $|\psi\rangle$, then $O|\psi\rangle = |\psi\rangle$.

[48] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2011), pp. 454–459.

[49] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, The Trotter step size required for accurate quantum simulation of quantum chemistry, *Quantum Inf. Comput.* **15**, 361 (2015).

[50] R. A. Lang, I. G. Ryabinkin, and A. F. Izmaylov, Unitary transformation of the electronic hamiltonian with an exact quadratic truncation of the Baker-Campbell-Hausdorff expansion, *J. Chem. Theory Comput.* **17**, 66 (2021).

[51] J. Dehaene and B. De Moor, Clifford group, stabilizer states, and linear and quadratic operations over GF(2), *Phys. Rev. A* **68**, 042318 (2003).

[52] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quantum Inf. Comput.* **12**, 901 (2012).

[53] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).

[54] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Exponential improvement in precision for simulating sparse Hamiltonians, in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2014), pp. 283–292.

[55] L. K. Grover, Quantum Mechanics Helps in Searching for a Needle in a Haystack, *Phys. Rev. Lett.* **79**, 325 (1997).

- [56] G. G. Guerreschi, Repeat-until-success circuits with fixed-point oblivious amplitude amplification, *Phys. Rev. A* **99**, 022306 (2019).
- [57] M. Boyer, G. Brassard, P. Hoyer, and A. Tapp, Tight bounds on quantum searching, *Fortschr. Physik* **46**, 493 (1998).
- [58] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, Tapering off qubits to simulate fermionic Hamiltonians, [arXiv:1701.08213](https://arxiv.org/abs/1701.08213).
- [59] W. M. Kirby, ContextualSubspaceVQE, 2021, <https://github.com/wmkirby1/ContextualSubspaceVQE>.
- [60] A. A. Hagberg, D. A. Schult, and P. J. Swart, Exploring network structure, dynamics, and function using networkx, in *Proceedings of the 7th Python in Science Conference*, edited by G. Varoquaux, T. Vaught, and J. Millman (SciPy, Pasadena, CA, 2008), pp. 11–15.
- [61] D. J. Welsh and M. B. Powell, An upper bound for the chromatic number of a graph and its application to timetabling problems, *Comput. J.* **10**, 85 (1967).
- [62] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [63] N. C. Rubin, R. Babbush, and J. McClean, Application of fermionic marginal constraints to hybrid quantum algorithms, *New J. Phys.* **20**, 053020 (2018).
- [64] J. F. Gonthier, M. D. Radin, C. Buda, E. J. Doskocil, C. M. Abuan, and J. Romero, Measurements as a roadblock to near-term practical quantum advantage in chemistry: resource analysis, *Phys. Rev. Res.* **4**, 033154 (2022).
- [65] A. Ralli and T. J. Weaving, symmer, 2022, <https://github.com/UCL-CCS/symmer>.