

Resource-Optimized Fermionic Local-Hamiltonian Simulation on a Quantum Computer for Quantum Chemistry

Qingfeng Wang¹, Ming Li², Christopher Monroe^{2,3}, and Yunseong Nam^{2,4}

¹Chemical Physics Program and Institute for Physical Science and Technology, University of Maryland, College Park, MD 20742, USA

²IonQ, College Park, MD 20740, USA

³Joint Quantum Institute, Department of Physics, and Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD 20742, USA

⁴Department of Physics, University of Maryland, College Park, MD 20742, USA

The ability to simulate a fermionic system on a quantum computer is expected to revolutionize chemical engineering, materials design, nuclear physics, to name a few. Thus, optimizing the simulation circuits is of significance in harnessing the power of quantum computers. Here, we address this problem in two aspects. In the fault-tolerant regime, we optimize the R_z and T gate counts along with the ancilla qubit counts required, assuming the use of a product-formula algorithm for implementation. We obtain a savings ratio of two in the gate counts and a savings ratio of eleven in the number of ancilla qubits required over the state of the art. In the pre-fault tolerant regime, we optimize the two-qubit gate counts, assuming the use of the variational quantum eigensolver (VQE) approach. Specific to the latter, we present a framework that enables bootstrapping the VQE progression towards the convergence of the ground-state energy of the fermionic system. This framework, based on perturbation theory, is capable of improving the energy estimate at each cycle of the VQE progression, by about a factor of three closer to the known ground-state energy compared to the standard VQE approach in the test-bed, classically-accessible system of the water molecule. The improved energy estimate in turn results in a commensurate level of savings of quantum resources, such as the number of qubits and quantum gates, required to be within a pre-specified tolerance from the known ground-state energy. We also explore a suite of generalized transformations of fermion to qubit operators and show that resource-requirement savings of up to more than 20%, in small instances, is possible.

1 Introduction

Simulating fermionic matter on a quantum computer has recently been receiving much attention. Already available in the literature are various chemistry and nuclear physics

Qingfeng Wang: keewang@umd.edu

Ming Li: li@ionq.co

Christopher Monroe: monroe@umd.edu

Yunseong Nam: nam@ionq.co

simulation results [8, 15, 19, 31, 34, 42], performed across multiple quantum computing platforms, including superconducting [8, 19, 34] and trapped-ion [15, 31, 42] based approaches. The spotlight on the simulation of fermionic systems on a quantum computer is not accidental. Simulations of these systems are useful for furthering fundamental science and practical engineering [13, 26], and quantum computers are expected to enable the quantum simulations of fermions undergoing local interactions [9, 23], a task believed to be classically difficult to scale.

Broadly speaking, simulations of fermionic systems on a quantum computer may be classified into two categories: a variational, quantum-classical hybrid simulation [37], suitable for imperfect, pre-fault tolerant (pre-FT) quantum computers, and a Hamiltonian dynamics simulation based on pure quantum simulation algorithms [6], typically considered in fault-tolerant (FT) quantum computers. In the context of estimating the ground-state energy of a fermionic system, the former leverages efficient preparation of ansatz states and evaluation of operator expectation values, both enabled by quantum computers. The latter leverages the ability of a quantum computer to efficiently simulate evolution of quantum systems with a local Hamiltonian, which, when combined with quantum phase estimation [47], allows us to evaluate the ground-state energy of the system.

In the pre-FT regime, quantum computational cost is dominated by the use of multi-qubit gates. In the FT regime, where quantum circuits are typically composed of gates in the Clifford+T gateset, quantum computational cost is dominated by the use of T := $\begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$ gates, many of which are used in the FT implementation of $R_z(\theta) := \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}$ gates (see, e.g., [4]). In this paper, we present approaches that optimize quantum simulations of fermionic systems in either the pre-FT regime, reducing the number of multi-qubit gates, or in the FT regime, reducing the number of R_z gates or the number of time steps that R_z gates need to be applied, in the simulation circuit.

Our paper is structured as follows. In Sec. 2 we briefly define the scope of the problem in both the pre-FT and the FT regimes and lay out the preliminaries. In Sec. 3, we present a highly-optimized circuit construction for the FT quantum simulations, where we improve R_z and T counts, as well as the number of ancilla qubits required over the state of the art. In Sec. 4, we show two complementary approaches that result in quantum resource savings in the pre-FT regime, i.e., a second-order perturbation based approach and a generalized fermion to qubit operator transformation. In the specific sample case of a water molecule, our method provides nearly 25% improvement over that obtained via the method detailed in [31] in the number of two-qubit gates required to achieve the accuracy to within chemical accuracy. We discuss our results in Sec. 5 and conclude in Sec. 6.

2 Preliminaries

We consider in this paper, as a concrete example, a fermionic system evolving according to a time-independent, local, second-quantized Hamiltonian H in the occupation basis

$$H = \sum_{p,r} h_{pr} a_p^\dagger a_r + \sum_{p,q,r,s} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s, \quad (1)$$

where a_p^\dagger and a_r denote the fermionic creation and annihilation operators on the p th and r th levels, respectively, and h_{pr} and h_{pqrs} denote single- and double- fermion Hamiltonian coefficients, respectively. The fermionic operators follow the canonical anti-commutation relations

$$\{a_j, a_k\} = 0, \quad \{a_j^\dagger, a_k^\dagger\} = 0, \quad \{a_j, a_k^\dagger\} = \delta_{jk} \mathbf{1}, \quad (2)$$

where $\{A, B\}$ denotes the anti-commutator $AB + BA$, δ_{jk} is the Kronecker delta, and $\mathbf{1}$ is the identity operator. For their implementations on a quantum computer, the fermionic operators need to be suitably transformed [5, 18, 43]. A well-known, popular choice is the Jordan-Wigner (JW) transformation [18], defined for an n -qubit system according to

$$a_j^\dagger = \mathbf{1}^{\otimes n-j-1} \otimes \sigma_+ \otimes \sigma_z^{\otimes j}, \quad a_j = \mathbf{1}^{\otimes n-j-1} \otimes \sigma_- \otimes \sigma_z^{\otimes j}, \quad (3)$$

where $j \in [0, n-1]$, \otimes denotes the tensor product, $\sigma_+ := \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$, $\sigma_- := \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$, and $\sigma_z := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$.

In the case where we consider a quantum dynamics simulation approach more suitable for the FT regime, we aim to implement the evolution operator $U_{\text{evo}} = e^{-iHt}$ on a quantum computer, where H is the system Hamiltonian and t is the duration by which we desire to evolve the system forward in time. Roughly, once the initial state $|\psi_{\text{init}}\rangle$, sufficiently close to the ground state $|\psi_{\text{ground}}\rangle$, is evolved to $|\psi(t)\rangle \approx e^{-iE_{\text{ground}}t}|\psi_{\text{init}}\rangle$ up to the closeness, quantum Fourier transform may be used to estimate the phase angle $E_{\text{ground}}t$, thus extracting the ground-state energy E_{ground} . Detailed discussion on the circuit layout that implements the quantum phase estimation algorithm is available in [33]. Note, an efficient implementation of the quantum Fourier transform on a FT quantum computer is known [32]. Therefore, the problem of an efficient energy-spectra calculation on a quantum computer in the FT regime boils down to the problem of efficiently implementing real-time quantum dynamics simulations.

A host of algorithms that (approximately) implement U_{evo} have been proposed, such as the asymptotically optimal quantum signal processing [24, 25]. However, it has been shown that for certain Hamiltonians including ones of the form (1) a more straightforward technique of the $2k$ th order product formula (PF) can be implemented more efficiently in practice [3, 6]. Thus in this paper we choose the PF algorithm for the FT regime quantum simulation and present methods to reduce the quantum resources required in its quantum circuit construction, measured in R_z gate counts and depths, in addition to the number of T gates and the ancilla qubits required (see Sec. 3).

In the case where we consider a quantum-classical hybrid approach more suitable for the pre-FT regime, we consider the widely-adopted technique of the variational quantum eigensolver (VQE) [37]. Specifically, we aim to implement a unitary ansatz evolution operator U_{ansatz} on a quantum computer, an example of which is the well-established unitary coupled cluster (UCC) ansatz [2, 16]. By tuning the variational parameters in U_{ansatz} , in a typical VQE approach, $\langle \psi_0 | U_{\text{ansatz}}^\dagger H U_{\text{ansatz}} | \psi_0 \rangle$ is minimized, where $|\psi_0\rangle$ is an initial state that is assumed to be close to the target ground state of H and can readily be prepared on a quantum computer. The goal of this hybrid approach is then to estimate the ground state energy of the fermionic system with the Hamiltonian of the form (1). We show in this paper a complementary approach, based on perturbation theory, to the hybrid method described above in order to better and more efficiently estimate the ground state energy (see Sec. 4.1). Our perturbative approach induces a negligible quantum resource overhead, measured in the number of multi-qubit gates. We also consider different fermion to qubit transformations other than the aforementioned JW transformation (see Sec. 4.2), which help reduce the number of multi-qubit gates without any loss in the algorithmic accuracy.

3 Fault-Tolerant Regime - Time evolution simulation

As discussed in Sec. 2, in this section, we detail the methods to optimize the time-evolution operator implementation on a quantum computer. In particular, we consider the PF

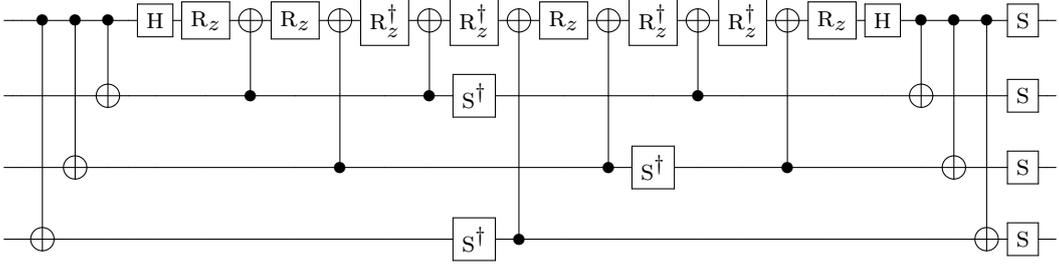


Figure 1: Standard two-body interaction circuit that implements $\exp[-i(\theta/2)\sigma_+\otimes\sigma_+\otimes\sigma_-\otimes\sigma_- + \text{h.c.}]$. Following the steps detailed in [31] closely, expanding $\sigma_+\sigma_+\sigma_-\sigma_- + \text{h.c.}$ (we suppress \otimes hereafter whenever contextually clear) into the particular ordering of $\sigma_x\sigma_x\sigma_x\sigma_x$, $\sigma_x\sigma_x\sigma_y\sigma_y$, $\sigma_x\sigma_y\sigma_y\sigma_x$, $\sigma_x\sigma_y\sigma_x\sigma_y$, $\sigma_y\sigma_y\sigma_x\sigma_x$, $\sigma_y\sigma_x\sigma_x\sigma_y$, $\sigma_y\sigma_x\sigma_y\sigma_x$, and $\sigma_y\sigma_y\sigma_y\sigma_y$, and implementing them one after another with the last qubit as the target qubit, we obtain the circuit shown in this figure after applying the circuit optimization routines detailed in [30, 31].

approach. The cost functions we consider are the number of R_z gates and the R_z gate depth, which are expected to be good proxy measures for the FT-regime quantum resource requirements. This is so since the de-facto standard gate set for the FT regime contains CNOT, single-qubit Clifford, and T gates, where T gates are assumed to be expensive as each implementation of the gate typically requires expending a carefully induced, distilled quantum state that is expensive to produce, and an R_z gate consumes multiple T gates for its precise implementation.

For completeness, we start by introducing the $2k$ th order PF algorithm, simulating the fermionic system described in Section 2. Assuming the JW transformation has been performed on the Hamiltonian H in (1), the time evolution operator we aim to implement may be written as

$$\exp\left(-i\sum_{j=1}^L\theta_j\hat{\sigma}^{(j)}\right) \approx [S_{2k}(\lambda)]^r, \quad (4)$$

where $\lambda := 1/r$, $\hat{\sigma}^{(j)} = \otimes_i \tilde{\sigma}^{(i,j)} + \text{h.c.}$, where $\tilde{\sigma}^{(i,j)} \in \{\sigma_+, \sigma_-, \sigma_z\}$ and h.c. denotes the Hermitian conjugate operator, and

$$\begin{aligned} S_1(\lambda) &:= \prod_{j=1}^L \exp(-i\theta_j\hat{\sigma}^{(j)}\lambda), \\ S_2(\lambda) &:= \prod_{j=1}^L \exp(-i\theta_j\hat{\sigma}^{(j)}\lambda/2) \prod_{j=L}^1 \exp(-i\theta_j\hat{\sigma}^{(j)}\lambda/2), \\ S_{2k}(\lambda) &:= [S_{2k-2}(p_k\lambda)]^2 S_{2k-2}((1-4p_k)\lambda) [S_{2k-2}(p_k\lambda)]^2, \end{aligned} \quad (5)$$

with $p_k := 1/(4 - 4^{1/(2k-1)})$ for $k > 1$ [44]. Inspecting (4) and (5), we observe that the individual exponential terms, hereafter referred to as a Trotter term, are of the form $\exp[-i\theta'_j(\otimes_i \tilde{\sigma}^{(i,j)} + \text{h.c.})]$, where θ'_j is a suitably scaled θ_j . A standard circuit that implements a Trotter term is readily available in [47]. Employing the circuit tricks in [31], in Fig. 1, we show an optimized quantum circuit that implements an example Trotter term $e^{-i\theta/2(\sigma_+\sigma_+\sigma_-\sigma_- + \text{h.c.})}$.

To more efficiently implement the two-body term in terms of R_z counts, we first directly evaluate the operator. It is straightforward to see that the two-body term $e^{-i\theta/2(\sigma_+\sigma_+\sigma_-\sigma_- + \text{h.c.})}$ implements an R_x -rotation between $|0011\rangle$ state and $|1100\rangle$ state. Note that, for an input

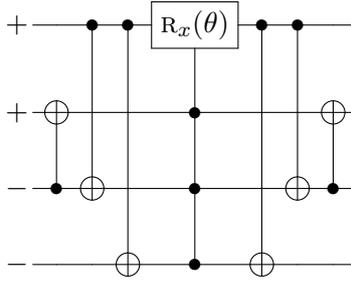


Figure 2: FT-regime optimized circuit for the two-body term $e^{-i\theta/2(\sigma_+\sigma_+\sigma_-\sigma_-+\text{h.c.})}$. We marked each qubit lines with either + or - to denote which qubits are associated with σ_+ or σ_- , respectively, for this particular example.

sate of $|b_0b_1b_2b_3\rangle$, where the m th qubit variable b_m is either 0 or 1, $|b_0 \ b_1 \oplus b_2 \ b_0 \oplus b_2 \ b_0 \oplus b_3\rangle$ maps $|0011\rangle$ to $|0111\rangle$ and $|1100\rangle$ to $|1111\rangle$. A triply-controlled R_x rotation on the zeroth qubit would thus implement the desired rotation on the mapped state, and all that needs to be done at this point is to map the rotated state back. We show a quantum circuit that implements the two-body term, constructed according to the aforementioned method, in Fig. 2.

We refer to Supplementary Material (SM) Sec. A for the circuit-level implementation details of Fig. 2. Briefly, a triply-controlled R_x gate in Fig. 2 may be implemented with Hadamard, R_z , and triply-controlled NOT gates. Specifically, we choose to use relative-phase triply-controlled NOT gates (see [27] and SM Sec. A for its decomposition into single- and two-qubit gates), as it admits less expensive circuit construction. Overall, our circuit decreases the R_z count required per two-body Trotter term to two.

Our construction can be contrasted to [39] wherein (see Fig. 6 of the appendix of [39]) an optimized implementation of a generic two-body term for the FT regime is discussed. Because the construction in [39] considers $\sigma_{x/y}$ -decomposed form of a two-body term with distinct coefficients – note our two-body term $e^{-i\theta/2(\sigma_+\sigma_+\sigma_-\sigma_-+\text{h.c.})}$ will have the same coefficients upon decomposition – the circuit construction there requires eight parallel R_z gates with different angles of rotation that are commensurate to the different coefficients. Considering our specific example, these eight angles would be made the same, and the parallel application of the same angle R_z gates is known to admit a more efficient implementation using the weightsum trick [10, 29]. As detailed in SM Sec. A, employing this trick reduces the gate requirement from the original $8R_z$ gates to $32T + 4R_z$ gates. The total number of ancilla qubits required is eleven in this approach. These may be compared with $16T + 2R_z$ gates and one ancilla qubit – note that a relative-phase triply-controlled NOT does not require an ancilla, unlike a conventional triply-controlled NOT gate (see SM Sec. A) – required for our method that is optimized for R_z depth. Both of these methods admit an R_z depth of one.

We note in passing that our method described above can be viewed as a generalization of the single-body term implementation discussed elsewhere [17, 21]. Indeed, the generalization may be applied to an arbitrary many-body term, should such a term be of interest to quantum simulations of a more complex kind. See SM Sec. A for detail for the single-body term description, inserted for completeness.

4 Pre-Fault Tolerant Regime - Variational Quantum Eigensolver

In the pre-FT regime, we aim to calculate the ground state energy of the system whose Hamiltonian is given in the form (1). This is typically achieved by the VQE approach. In this approach, by iteratively calling the quantum computer to compute the energy of parametrized ansatz states, one aims to minimize the energy and variationally obtain the ground state of the target system.

We consider in this paper a type of ansatz states $|\Psi_{\text{ansatz}}\rangle$ that are transformed from an initial state $|\Psi_0\rangle$ using a parametrizable unitary ansatz evolution operator U_{ansatz} , as $|\Psi_{\text{ansatz}}\rangle = U_{\text{ansatz}}|\Psi_0\rangle$. For brevity, we drop the subscript in U_{ansatz} in the following text. Since U is a unitary operator, it can always be written in the exponential form e^{Z-Z^\dagger} and Z is in turn parametrized. By varying the parameters in Z , the ansatz state energy $\langle\Psi_{\text{ansatz}}|H|\Psi_{\text{ansatz}}\rangle = \langle\Psi_0|U^\dagger H U|\Psi_0\rangle$ is variationally minimized. With a proper fermion to qubit basis transformation the energy expectation value can be evaluated efficiently on a quantum computer, where the quantum resource cost of this hybrid approach is the implementation cost of the operator U . We note that the procedures proposed in this section can be implemented with any ansatz with a unitary evolution operator U . As a concrete example for discussion, we use the widely-adopted UCC ansatz with single and double excitations (UCCSD), a systematic method that is universally applicable to any quantum hardware backend. Other ansatzes, for instance the k -UpCCGSD ansatz of Lee et al. [22] or the UCCGSD ansatz of Grimsley et al. [12], can also be readily applied.

The UCCSD method starts with a ground state Hartree-Fock (HF) wavefunction $|\Psi_0\rangle$ that can be easily calculated on a classical computer and implemented on a quantum computer. It then evolves $|\Psi_0\rangle$ with a unitary operator $U = e^{Z_{\text{UCCSD}} - Z_{\text{UCCSD}}^\dagger}$, where $Z_{\text{UCCSD}} = \sum_{l=1}^2 Z_l$ is the so-called cluster operator. $Z_{1,2}$ are the single and double excitation operators and are written in second quantization as

$$\begin{aligned} Z_1 &= \sum_{\substack{p \in \text{virt} \\ r \in \text{occ}}} t_{pr} a_p^\dagger a_r, \\ Z_2 &= \sum_{\substack{p, q \in \text{virt} \\ r, s \in \text{occ}}} t_{pqrs} a_p^\dagger a_q^\dagger a_r a_s, \end{aligned} \tag{6}$$

where t_{pr} and t_{pqrs} are variational parameters and ‘‘virt’’ and ‘‘occ’’ denote virtual and occupied levels respectively. Notice that for the UCCSD ansatz, to minimize the size of the circuit that implements U , not all terms in the cluster operator are needed to achieve a certain precision. In other words, only a subset of all the parameters t_{pq} and t_{pqrs} need to be included in the VQE approach, while the rest can be set to zero. In fact, it is very much the name of the game in the pre-FT regime to find a minimal set of variational parameters so that the final ground state energy satisfies certain error criteria.

In this section, we propose two general procedures to reduce circuit complexity of the approach detailed above, represented by the total number of multi-qubit gates. To be concrete with our example, we use the first-order PF algorithm to implement the UCCSD ansatz, although extensions to other ansatzes or higher order PF algorithms are straightforward. In Sec. 4.1, we detail a hybrid framework that is based on perturbation theory, wherein we apply a perturbation correction to each instance of the VQE circuit. As will be discussed in Sec. 4.1 in detail, our perturbation approach naturally provides a means to determine the variational parameters to further include in the ansatz for a given VQE circuit. This enables us to systematically build, in an iterative fashion, a larger and larger VQE circuit, effectively bootstrapping the VQE progression towards the ground-state energy

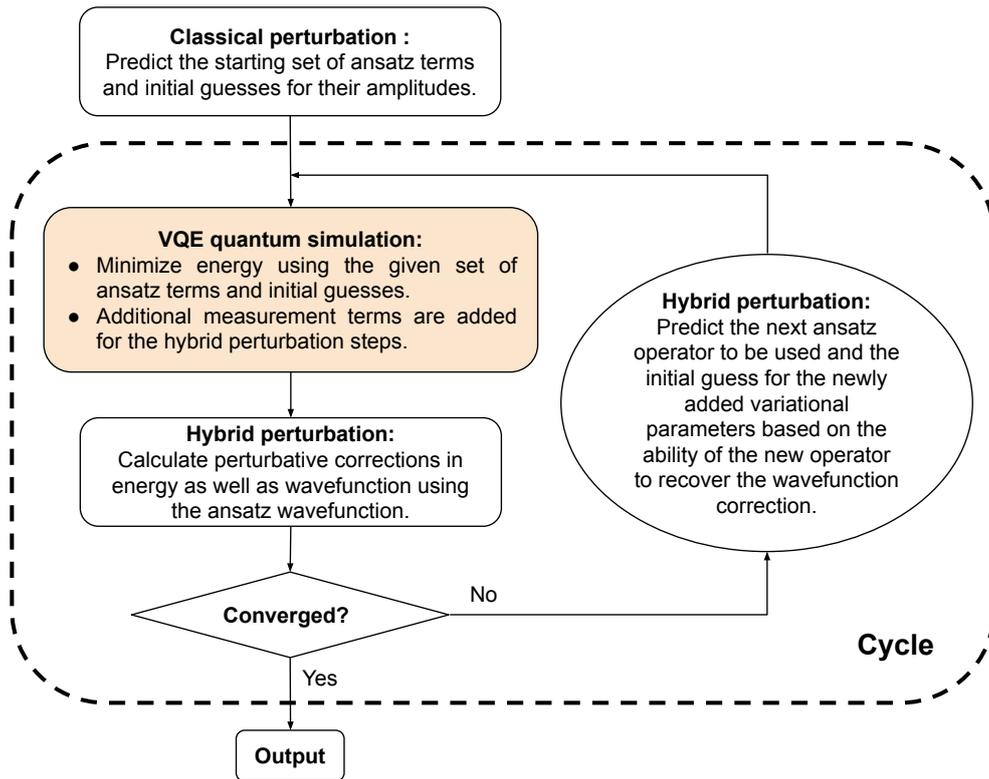


Figure 3: Flow diagram for the proposed framework incorporating perturbation methods to VQE simulations. See main text for detailed description.

estimate of the simulated system. In Sec. 4.2, we explore a suite of generalized fermion to qubit operator transformations to reduce the cost of quantum computation without sacrificing algorithmic accuracy.

4.1 Perturbation Assisted Quantum Simulation

In this section, we describe a general framework that leverages the power of perturbation theory to optimize VQE-based quantum simulations by predicting the subset of variational parameters to include in the ansatz and correcting the VQE result via post processing. The framework optimizes both the total number of VQE executions as well as the size of the ansatz state preparation circuits used to reach convergence in the ground-state energy estimate. In Sec. 4.1.1 we outline the framework. The derivation of a simple perturbation scheme that can be straightforwardly implemented in the framework is shown in Sec. 4.1.2, which we hereafter refer to as a hybrid second order Møller-Plesset perturbation (HMP2) method. In Sec. 4.1.3 we present a classically simulated comparison between the HMP2 scheme and a more conventional VQE approach [31], by computing the ground state energy of a water molecule using the UCCSD ansatz.

4.1.1 Perturbative predictor and corrector

We detail, in this section, a general, systematic framework of using perturbation methods to improve quantum resources required for the VQE approach. Specifically, we aim to rapidly converge to the ground state energy with small quantum circuits.

Any general (not specific to any system) unitary ansatz operator can be written as

$$U = e^{Z-Z^\dagger}, \quad Z = \sum_{\{\alpha\}} f_\alpha(\{t_\beta\}) D_\alpha, \quad (7)$$

where the summation is over a set of orbital sequences $\{\alpha\}$, f_α are real functions of the set of variational parameters $\{t_\beta\}$, and D_α is a generic orbital substitution operator that substitutes the orbitals in a wavefunction according to the orbital sequence α . For instance, a single orbital substitution operator $D_{p^\dagger q} = a_p^\dagger a_q$ substitutes orbital q with p . We can then define the orbital substituted wavefunction $|\Psi_{D_\alpha}\rangle = D_\alpha|\Psi_0\rangle$. For a specific general ansatz approach there is a full set of variational parameters $T_{\text{full}} = \{t_\beta\}$ that can be used to parametrize the ansatz evolution operator. In the case of UCCSD, by comparing (6) with (7), we note that $T_{\text{full}} = \{t_{pr}, t_{pqrs} \mid p, q \in \text{virt}; r, s \in \text{occ}\}$, $\{\alpha\} = \{p^\dagger r, p^\dagger q^\dagger rs \mid p, q \in \text{virt}; r, s \in \text{occ}\}$, and $f_\alpha = t_\alpha$.

Note further that we can generate a set of ansatz evolution operators $\{U_{T_i}\}$, each element of which is the unitary operator that induces an ansatz state with a subset T_i of the full variational parameters T_{full} ($T_i \subset T_{\text{full}}$). The rest of the parameters in the complement set $T_{\text{full}} \setminus T_i$ are fixed to an ordered set of default values $T_{\text{full}}^{(0)}$. Without loss of generality, we assume the default values are absorbed into f_α in (7) and the elements in the complement set $T_{\text{full}} \setminus T_i$ are simply set to zero. For every U_{T_i} in $\{U_{T_i}\}$, the ansatz state energy can be variationally minimized as $E_{T_i} = \mathbf{min}_{\text{elements}(T_i) \in \mathbb{R}} \langle \Psi_0 | U_{T_i}^\dagger H U_{T_i} | \Psi_0 \rangle$, where $\text{elements}(T_i)$ denotes the elements of the set T_i . For an application of a particular general ansatz through the VQE approach in the pre-FT regime, one aims to find a U_{T_i} that has the least circuit complexity while satisfying $|E_{T_i} - E_{T_{\text{full}}}| < \varepsilon$ with ε a pre-determined error criterion.

A concrete strategy to reach convergence with minimal circuit complexity is to start with an operator with a small set of variational parameters that is known to be necessary (but not necessarily sufficient) to achieve the error threshold and iteratively grow it. Going from the m th iteration to the $(m+1)$ th one, the set of variational parameters in the operators grows and satisfies $T_m \subset T_{m+1}$. A critical part in ensuring a rapid convergence is the selection of the additional variational parameter(s) to include between iterations. In the case of UCCSD, they directly correspond to additional individual excitation terms in (6) used to prepare the ansatz state. Previous works have, e.g., used the full configuration interaction (FCI) results [31] in the case the system is sufficiently small to be simulatable on a classical computer, and have demonstrated the significance of the excitation term selection for the resource requirement. Here, we propose to iteratively select the next variational parameters to include in the ansatz state preparation based on the size of the perturbatively predicted wavefunction correction amplitude for a given ansatz state whose variational parameters are already optimized via the conventional means of VQE. Our strategy also evaluates a perturbative energy correction in addition to the conventional VQE approach, directly contributing to the fast convergence, while incurring no overhead in the pre-FT regime resource requirement.

Figure 3 shows the flow diagram of the proposed framework. As in the conventional VQE approach for fermionic simulations, we start with the ground state Hartree-Fock wavefunction of the single particle Hamiltonian as $|\Psi_0\rangle$. The first step is to determine the initial evolution operator U_{T_1} to use for the first iteration of VQE. As an example, using the two-particle Hamiltonian as perturbation, energy and wavefunction corrections can be calculated using classical algorithms with relative ease. From the perturbed wavefunction, we can extract the amplitudes of individual excited states in the single-particle Hamiltonian basis. These amplitudes will serve as the initial guesses of the variational parameters for the

first round of VQE simulation, which could significantly reduce the number of evaluations of the quantum circuit comparing to all-zeros or random initial guesses [40]. If we demand the energy convergence criteria of δ_E – The entire simulation is considered to be converged when the magnitude of energy change associated with an addition of any extra ansatz term is smaller than δ_E – we include in the initial ansatz set the ansatz terms whose contribution to the correlation energy is greater than $f(\delta_E)$, where a standard choice of function $f(\delta_E)$ may simply be δ_E .

With the initial set of ansatz terms and their initial variational parameter values, we run the first round of VQE simulation to minimize the energy $E_{T_1} = \langle \Psi_0 | U_{T_1}^\dagger H U_{T_1} | \Psi_0 \rangle$. Once the energy is minimized and the ansatz state converges, we proceed to compute the expectation values of a set of operators that were chosen in advance to inform us about the perturbation around the converged ansatz state. We refer to this method as a hybrid perturbation since unlike the conventional perturbation where we know the unperturbed wavefunction in advance, we start the perturbation from the converged VQE ansatz state.

Based on the additional measurements performed on a quantum computer, we may now use the hybrid perturbation method to calculate corrections to the correlation energy and the wavefunction. The resulting total energy of this cycle is the summation of the energy correction and the VQE energy. Before starting the next cycle, we first identify a pool of operators that can be the immediate successor of the current operator. Then we evaluate how much each operator in the pool can account for towards the perturbative wavefunction correction. The operator that can recover the wavefunction correction the most gets the nod for the next cycle and the way it recovers the correction informs the initial guesses of the newly added variational parameters in the next cycle. A more detailed illustration of such a procedure is given in the next subsection. The initial variational parameter values of the ansatz terms that continue to be a part of the next VQE simulation may simply be imported from the cycle before.

The cycle iterates the outlined procedure of running the VQE simulation and the hybrid perturbation calculation, until convergence of the total energy is achieved. We next detail the steps of a simple implementation of our hybrid perturbation framework that is inspired by the classical MP2 method.

4.1.2 HMP2 method

We derive in this section our hybrid perturbation method, inspired by the Møller-Plesset perturbation theory, applied to the VQE simulation. Our hybrid MP2-based perturbation method, HMP2, aims to help each VQE cycle in terms of both improving the resulting energy via a perturbative energy correction and optimizing the ansatz operator to use in the next cycle through a perturbative wavefunction correction.

In the m th VQE cycle, we can write the energy of the ansatz state E_{T_m} in the context of first-order perturbation theory, i.e.,

$$\begin{aligned}
 E_{T_m} &= \langle \Psi_0 | U_{T_m}^\dagger H U_{T_m} | \Psi_0 \rangle \\
 &= \langle \Psi_0 | F + (U_{T_m}^\dagger H U_{T_m} - F) | \Psi_0 \rangle \\
 &= \langle \Psi_0 | F | \Psi_0 \rangle + \langle \Psi_0 | U_{T_m}^\dagger H U_{T_m} - F | \Psi_0 \rangle \\
 &= E_0 + (E_{T_m} - E_0) \\
 &= E^{(0)} + E_{T_m}^{(1)},
 \end{aligned} \tag{8}$$

where F is the Fock operator, E_0 is the sum of orbital energies, $E^{(0)} = E_0$ is the zeroth-order energy, and $E_{T_m}^{(1)} = E_{T_m} - E_0$ is the first-order correction energy. Based on pertur-

bation theory, a second-order correction to the energy can now be written as

$$E_{T_m}^{(2)} = \sum_{\{\alpha'\}} \frac{|\langle \Psi_{D_{\alpha'}} | V_{T_m} | \Psi_0 \rangle|^2}{\Delta E_{D_{\alpha'}}}, \quad (9)$$

where $V_{T_m} = U_{T_m}^\dagger H U_{T_m} - F$. The set of orbital sequences $\{\alpha'\}$ does not necessarily need to be the same as the set $\{\alpha\}$ used in (7) and should be chosen according to accuracy requirement and resource constraints. See Sec. 5 for more detailed discussion. In our case study using the UCCSD ansatz, we choose $\{\alpha'\} = \{\alpha\}$. The energy $\Delta E_{D_{\alpha'}}$ is defined as the orbital energy difference of the orbital substitutions. For instance, $\Delta E_{D_{p \dagger q}} = E_q - E_p$, where E_p and E_q are the orbital energies of the p th and the q th orbitals, respectively. Inserting $V_{T_m} = U_{T_m}^\dagger H U_{T_m} - F$ in (9), the numerator becomes

$$\begin{aligned} |\langle \Psi_{D_{\alpha'}} | V_{T_m} | \Psi_0 \rangle|^2 &= |\langle \Psi_{D_{\alpha'}} | U_{T_m}^\dagger H U_{T_m} - F | \Psi_0 \rangle|^2 \\ &= \left| \langle \Psi_{D_{\alpha'}} | U_{T_m}^\dagger H U_{T_m} | \Psi_0 \rangle - \langle \Psi_{D_{\alpha'}} | F | \Psi_0 \rangle \right|^2 \\ &= \left| \langle \Psi_{D_{\alpha'}} | U_{T_m}^\dagger H U_{T_m} | \Psi_0 \rangle \right|^2, \end{aligned} \quad (10)$$

where we used $\langle \Psi_{D_{\alpha'}} | F | \Psi_0 \rangle = 0$.

In order to apply the VQE results, especially the ansatz wavefunction $|\Psi_{\text{ansatz}}^{T_m}\rangle$ directly to the perturbation calculation, we proceed as follows. For brevity, we introduce $\tilde{Z}_{T_m} = Z_{T_m} - Z_{T_m}^\dagger$ with $U_{T_m} = e^{\tilde{Z}_{T_m}}$ and $\tilde{D}_{\alpha'} = D_{\alpha'} - D_{\alpha'}^\dagger$. Inserting these into (10), we obtain

$$\begin{aligned} |\langle \Psi_{D_{\alpha'}} | U_{T_m}^\dagger H U_{T_m} | \Psi_0 \rangle|^2 &= |\langle \Psi_0 | \tilde{D}_{\alpha'}^\dagger U_{T_m}^\dagger H U_{T_m} | \Psi_0 \rangle|^2 \\ &= \left| \langle \Psi_{\text{ansatz}}^{T_m} | U_{T_m} \tilde{D}_{\alpha'}^\dagger U_{T_m}^\dagger H | \Psi_{\text{ansatz}}^{T_m} \rangle \right|^2. \end{aligned} \quad (11)$$

Next we Taylor expand the $e^{\pm \tilde{Z}_{T_m}}$ in $e^{\tilde{Z}_{T_m}} \tilde{D}_{\alpha'}^\dagger e^{-\tilde{Z}_{T_m}} H$ up to first order in \tilde{Z}_{T_m} . This is consistent with our choice of the first-order PF algorithm for the implementation of the ansatz. We obtain

$$\begin{aligned} & \left| \langle \Psi_{\text{ansatz}}^{T_m} | e^{\tilde{Z}_{T_m}} \tilde{D}_{\alpha'}^\dagger e^{-\tilde{Z}_{T_m}} H | \Psi_{\text{ansatz}}^{T_m} \rangle \right|^2 \\ &= \left| \langle \Psi_{\text{ansatz}}^{T_m} | (1 + \tilde{Z}_{T_m} + \dots) \tilde{D}_{\alpha'}^\dagger (1 - \tilde{Z}_{T_m} + \dots) H | \Psi_{\text{ansatz}}^{T_m} \rangle \right|^2 \\ &\approx \left| \langle \Psi_{\text{ansatz}}^{T_m} | \tilde{D}_{\alpha'}^\dagger H - \tilde{D}_{\alpha'}^\dagger \tilde{Z}_{T_m} H + \tilde{Z}_{T_m} \tilde{D}_{\alpha'}^\dagger H | \Psi_{\text{ansatz}}^{T_m} \rangle \right|^2. \end{aligned} \quad (12)$$

With (9) and (12) we can obtain the energy correction to the VQE result.

We note that $\tilde{D}_{\alpha'}^\dagger H - \tilde{D}_{\alpha'}^\dagger \tilde{Z}_{T_m} H + \tilde{Z}_{T_m} \tilde{D}_{\alpha'}^\dagger H$ is a sum of products of Pauli operators $\sum_j \epsilon_j \hat{\sigma}^{(j)}$ in the qubit basis, after applying a suitable fermion to qubit basis transformation. Since $\hat{\sigma}^{(j)}$ have eigenvalues $+1$ or -1 , denoting $|\Psi_j^{(p)}\rangle$ and $|\Psi_j^{(m)}\rangle$ as the eigenvectors with the respective eigenvalues, we may then write

$$\begin{aligned} & \left| \langle \Psi_{\text{ansatz}}^{T_m} | \sum_j \epsilon_j \hat{\sigma}^{(j)} | \Psi_{\text{ansatz}}^{T_m} \rangle \right|^2 \\ &= \left| \langle \Psi_{\text{ansatz}}^{T_m} | \sum_j \epsilon_j \left(\sum_p |\Psi_j^{(p)}\rangle \langle \Psi_j^{(p)}| - \sum_m |\Psi_j^{(m)}\rangle \langle \Psi_j^{(m)}| \right) | \Psi_{\text{ansatz}}^{T_m} \rangle \right|^2 \\ &= \left| \sum_j \epsilon_j \left(\sum_p |\langle \Psi_{\text{ansatz}}^{T_m} | \Psi_j^{(p)} \rangle|^2 - \sum_m |\langle \Psi_{\text{ansatz}}^{T_m} | \Psi_j^{(m)} \rangle|^2 \right) \right|^2. \end{aligned} \quad (13)$$

Note that (13) requires only a simple projection of $|\Psi_{\text{ansatz}}\rangle$ onto $|\Psi_j^{(p)}\rangle$ or $|\Psi_j^{(m)}\rangle$. Thus the second order correction energy may be obtained without any quantum resource overhead in the circuit size.

To optimize the choice of the ansatz evolution operator to use in the next cycle, we turn to the wavefunction correction. The ansatz state wavefunction can be considered as the zeroth-order wavefunction $|\Psi^{(0)}\rangle$ and a first order correction is given by

$$|\Psi_{T_m}^{(1)}\rangle = \sum_{\{\alpha\}} \frac{\langle \Psi_{D_\alpha} | V_{T_m} | \Psi_0 \rangle}{\Delta E_{D_\alpha}} |\Psi_{D_\alpha}\rangle. \quad (14)$$

We then use the following procedure to determine which ansatz evolution operator to use next in the $(m+1)$ th round. In the m th round we have variationally determined a set of values for the parameters in the set T_m , denoted as S_m . We then look for a set $R_m = \{U_{T_i}\}$ of all the evolution operators U_{T_i} that satisfies $T_m \subset T_i$ and $n(T_i \setminus T_m) = \min_{j \neq m} [n(T_j \setminus T_m)]$, where $n(A)$ denotes the number of elements in set A . R_m is then the pool of operators from which we are going to choose the next ansatz evolution operator. For each U_{T_i} in R_m , we can then compute the overlap $F_m^i = 1/w_m^i \sum_{t_\beta \in T_i \setminus T_m} |\langle \Psi_{t_\beta}^{m,i} | \Psi_{T_m}^{(1)} \rangle|$, where

$$|\Psi_{t_\beta}^{m,i}\rangle = \left(\frac{\partial U_{T_i}}{\partial t_\beta} \Big|_{S_m} \right) |\Psi_0\rangle. \quad (15)$$

Here we use w_m^i to capture the additional cost associated with implementing U_{T_i} instead of U_{T_m} . For example, it can be the number of additional two-qubit gates needed to compute $\langle \Psi_0 | U^\dagger H U | \Psi_0 \rangle$ for $U = U_{T_i}$ instead of $U = U_{T_m}$. For our concrete example, we simply take w_m^i to be the number of additional fermionic operators in T_i comparing to T_m . Using the identity $\langle \Psi_{D_\alpha} | \Psi_{D_{\alpha'}} \rangle = \delta_{\alpha\alpha'}$, it is straightforward to evaluate all of the F_m^i using (9), (12), and (15). Then we simply pick the U_{T_i} out of R_m with the largest F_m^i to be the ansatz evolution operator to use in the next cycle. We can also use the wavefunction correction to guess the initial values for an additional variational parameter $t_\beta \in T_{m+1} \setminus T_m$ as $\pm \langle \Psi_{t_\beta}^{m,m+1} | \Psi_{T_m}^{(1)} \rangle$.

In the case of UCCSD, each U_{T_i} in R_m only has one more D_α term in Z_{T_i} than Z_{T_m} of U_{T_m} , which corresponds to an incremental change in the terms included in the ansatz operator Z . Using $e^{\tilde{Z}_{T_i}} \approx 1 + \tilde{Z}_{T_i}$, F_m^i is approximately the perturbative amplitude $|\langle \Psi_{D_\alpha} | V_{T_m} | \Psi_0 \rangle / \Delta E_{D_\alpha}|$ of the corresponding additional D_α .

All the procedures introduced above for the m th round of a VQE cycle can also be used before the first cycle by simply taking $m = 0$. In this case, all the calculations are classical and the energy correction is simply the classical MP2 energy correction. The wavefunction correction however can inform the first operator to be tried and suggest the initial values for its variational parameters.

We note in passing that, in principle, the number of individual Pauli-string operators in (13) whose expectation values are to be evaluated using a quantum computer scales as $O(n^4)$ per UCCSD ansatz where n is the number of qubits. This is so since H consists of $O(n^4)$ terms, while, for a particular ansatz, Z is fixed and provides a constant prefactor. Thus the predictive feature using (14) requires $O(n^8)$ Pauli-string measurements because there are $O(n^4)$ elements in $\{\alpha\}$. The perturbative correction using (9) can reuse all the measurement results from the predictive feature because we choose $\{\alpha'\} = \{\alpha\}$. While this may appear challenging, in practice, a series of techniques can be applied to significantly reduce the number of evaluations. First, we can take advantage of the fact that only the qubits that are coupled via a chosen ansatz are entangled. By carefully choosing the qubits

to use to create the ansatz state, we can treat the qubits that have not been operated on classically. A straightforward extension to small, disjoint sets of qubits with set sizes that admit inexpensive classical postprocessing may be considered as well, although we do not consider such an aggressive optimization in this paper. In addition, if we have two qubits that represent the same spatial orbital with two opposite spins, in the case where they are not distinguishable in their energy due to a particular choice of the ansatz (see Sec. 4.2 and [31] for more details), the two qubits would encode a redundant piece of information, and this allows us to encode the information using just one qubit. We call this optimization technique qubit space reduction (QSR) (see SM Sec. H for the implementation details). Secondly, it is possible to measure multiple Pauli strings simultaneously to reduce the total number of measurements provided that these Pauli strings commute with each other. For instance, two methods, the general commuting (GC) partition and the qubit-wise commuting (QWC) partition, have been proposed to enable such simultaneous measurements [11]. Generally speaking, GC can reduce the number of measurements significantly but could incur additional cost in terms of additional two-qubit gates while QWC can reduce the number of measurements, though to a lesser degree than GC, without increasing the number of two-qubit gates. If using GC, it is beneficial to use QSR beforehand to reduce the size of the qubit space needed to be partitioned, which in turn reduces the number of extra two-qubit gates incurred from using the GC partition.

4.1.3 Comparison to prior state-of-the-art

To demonstrate our framework of perturbation assisted quantum simulation using the HMP2 method, we performed classically emulated VQE calculations with the UCCSD ansatz of the ground state energy of a water molecule at its equilibrium geometry. Using the STO-3G basis, the calculation contains 14 qubits in total. Table 1 shows the incremental changes of the UCCSD correlation energy as well as the HMP2 correction as more ansatz terms are included according to our framework. Note that our choice of $\{\alpha'\}$ in (9) for this example guarantees that the final energy reaches the UCCSD energy when all the ansatz terms are included in (6). To compare, we also listed the classically calculated HF, MP2, CCSD, and FCI energies since it is difficult to obtain the UCCSD energy via classical algorithms. With the HMP2 correction, the total energy $E_{\text{UCCSD+HMP2}}$ descends quickly towards FCI energy.

Figure 4 shows the convergence of the ground state energies using different approaches as the number N of D_α terms included in the UCCSD ansatz operator increases. For the conventional approach [31], the insertion order of the ansatz terms is obtained from the prior knowledge of the order of contribution of determinants in a classical FCI calculation, which closely mimics the ideal case but is rarely realistic to obtain. For the UCCSD energies obtained using the proposed HMP2, we bootstrapped the ordering of the ansatz terms as detailed in Sec. 4.1.2. Comparing the convergence of the UCCSD energies, we find that the HMP2-bootstrapped ordering effectively captures the major ansatz terms. This is confirmed by the good agreement between the FCI ordering and the HMP2 ordering shown in Fig. 4. We also observe that the HMP2 correction to the ground-state energy helps accelerate the energy convergence towards the FCI energy significantly. Provided that implementation of each additional ansatz term leads to a substantial accumulation of noise in the NISQ hardware, we find the rapid energy convergence enabled by the HMP2 method to be particularly useful for the near-term quantum computers.

In addition to focusing on reducing the circuit complexity, which is the most important cost for NISQ era applications, we examine the total number of measurements as a secondary cost using the water molecule calculation as an example. Figure 5 (a) shows

	Ansatz	E_{UCCSD}	$E_{\text{P}_D}^{\text{corr}}$	$E_{\text{UCCSD+HMP2}}$
Ansatz terms in HMP2 order	HF+1	-74.9749	-0.0249	-74.9998
	HF+2	-74.9781	-0.0220	-75.0001
	HF+4	-74.9854	-0.0170	-75.0024
	HF+5	-74.9881	-0.0155	-75.0036
	HF+6	-74.9909	-0.0139	-75.0048
	HF+8	-74.9966	-0.0091	-75.0057
	HF+9	-74.9996	-0.0068	-75.0063
	HF+11	-75.0038	-0.0039	-75.0077
	HF+12	-75.0047	-0.0034	-75.0081
	HF+14	-75.0074	-0.0023	-75.0098
	HF+16	-75.0091	-0.0015	-75.0106
	HF+17	-75.0100	-0.0009	-75.0109
	HF+19	-75.0102	-0.0007	-75.0109
	HF+21	-75.0107	-0.0004	-75.0111
	HF+23	-75.0111	-0.0002	-75.0113
	HF+24	-75.0113	-0.0001	-75.0114
	HF+26	-75.0113	-0.0001	-75.0114
	HF+28	-75.0113	-0.0001	-75.0114
	E_{HF}			-74.9624
	E_{MP2}			-74.9977
	E_{CCSD}			-75.0114
	E_{FCI}			-75.0116

Table 1: Ground state energy calculations for a water molecule using STO-3G basis. HF+ N represents a VQE cycle with the UCCSD ansatz with N D_α terms, as described in Sec. 4.1.1. $E_{\text{P}_D}^{\text{corr}}$ is the HMP2 correction based on (12). $E_{\text{UCCSD+HMP2}} = E_{\text{UCCSD}} + E_{\text{P}_D}^{\text{corr}}$ is the total energy obtained in one VQE cycle. For the classical optimization step in VQE, we used L-BFGS-B optimizer [48], and we assumed perfect measurement with an infinite number of circuit repetitions to suppress the effect of finite measurement. The classically computed energies E_{HF} , E_{MP2} , E_{CCSD} , and E_{FCI} for the water molecule with the same geometry and basis set are also listed for comparison. All energies are in units of Hartree.

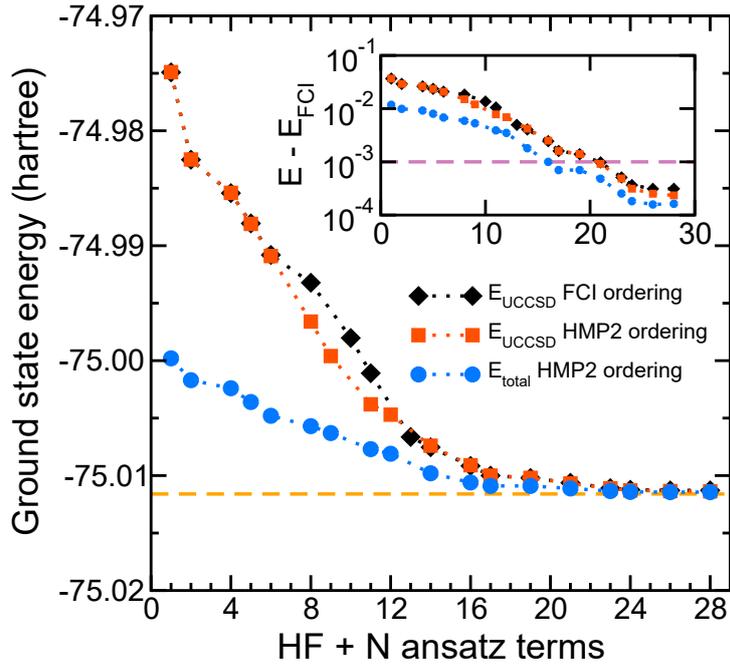


Figure 4: Comparison of the ground state energies of a water molecule at its equilibrium geometry using the STO-3G basis set, calculated by various methods as a function of N , the number of D_α terms included in the HF+N ansatz. The orange dashed line is the FCI energy calculated using the PSI4 package [35], which serves as the benchmark. The black diamonds connected by the dotted lines are the UCCSD energies E_{UCCSD} , calculated using different numbers of ansatz terms ordered by the contribution of corresponding determinants to the FCI energy. The red squares and blue circles connected by the dotted lines are the ground-state energies computed according to the proposed framework with the HMP2 ordering, with the circles containing the additional HMP2 energy correction at each VQE cycle. The inset shows in semi-log the differences between the energies obtained by the aforementioned methods and the FCI energy as a function of N . The purple dashed line shows the chemical accuracy given by 10^{-3} hartree. As stated in the caption of Table 1 as well, for the classical optimization step in VQE, we used L-BFGS-B optimizer [48], and we assumed perfect measurement with an infinite number of circuit repetitions to suppress the effect of finite measurement.

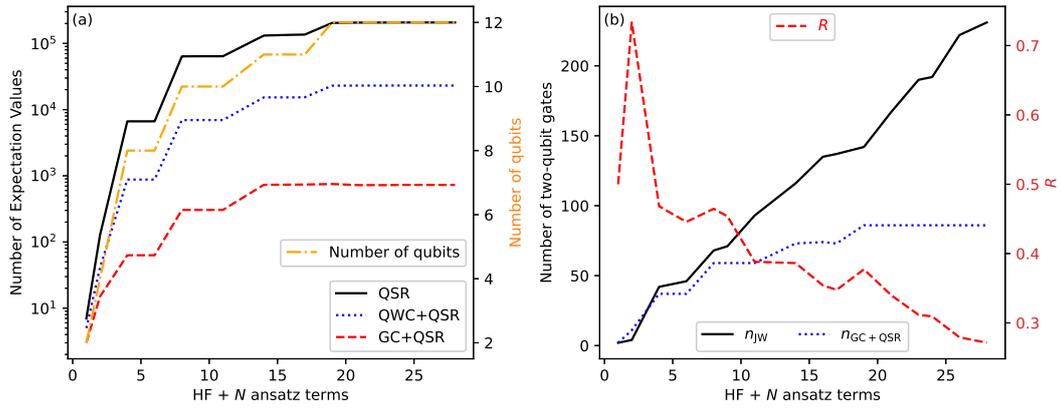


Figure 5: Demonstration of the reduction of the number of measurements as a secondary cost for the water molecule calculation using various optimization schemes. The calculation is carried out using the HMP2 framework with the UCCSD ansatz. The total number of expectation values to be evaluated on a quantum computer as a function of the number N of D_α terms in the UCCSD ansatz operator for a VQE cycle is shown in (a) for three different optimization schemes. The black solid line represents the reduction using only QSR. The blue dotted line shows the reduction using the QWC partition and QSR at the same time. The red dashed line illustrates the most reduction, using the GC partition along with QSR. Note the plateau-like behaviors that make the lines to appear like staircases are due to our QSR technique (orange, dot-dashed line with the y axis on the right), i.e., the jumps in the number of expectation values to be measured per VQE iteration occur when the QSR technique makes a transition from one space requirement to another. The average number of additional two-qubit gates incurred by using the GC partition along with QSR, n_{GC+QSR} , is shown in (b) as the blue dotted line. This may be compared with the black solid line, representing the total number of two-qubit gates n_{JW} required to induce UCCSD ansatz states, implemented with the JW transformation. The distribution of n_{GC+QSR} is given in the SM Sec. I. The red dashed line with an alternative y -axis label on the right shows the ratio $R = n_{GC+QSR} / (n_{JW} + n_{GC+QSR})$.

the number of measurements using various optimization schemes discussed in Sec.4.1.2. Comparing to the number of measurements optimized using only QSR, the number optimized using the GC partition method and QSR is more than two orders of magnitude smaller while the number is about one order of magnitude smaller using the QWC partition method along with QSR. Figure 5 (b) shows the additional cost in terms of the number of two-qubit gates associated with using the GC partition method along with QSR comparing to the total number of two-qubit gates using the JW transformation. The ratio between the two quickly reduces as the number of ansatz terms increases. Thus one should decide whether to use GC or QWC along with QSR after weighing their benefits against their cost on a case-by-case basis.

4.2 Generalized transformations for Fermion to Qubit operator

In this section, we investigate how a variety of fermion to qubit transformations may be used to reduce the quantum resource requirements for the pre-FT fermion simulations. It is important to note that all of the transformations are equivalent and thus the resulting quantum circuits for each transformation implements exactly the same fermion simulation. Therefore, the resource savings we obtain in this section are independent of the accuracy of the simulations, in general.

Well-known fermion to qubit transformations, such as Jordan-Wigner (JW) [18] or Bravyi-Kitaev (BK) [5] transformations, map fermionic creation (annihilation) operators to Pauli strings. However, there exist numerous other transformations available for use. Below, we introduce a generalized transformation (GT) method (see also [43]), of which the JW and BK transformations are a part. We show that, when used with the PF algorithm as a concrete example for the implementation of the UCCSD ansatz, significant quantum resource savings may be achieved by a suitable choice of the mapping for a given cluster operator input, together with a carefully chosen sequence of heuristic optimization methods.

All transformations in the GT method must respect relations specified in (2). This may be achieved by considering the following invertible, upper-triangular basis-transformation matrix β , which transforms the occupation number basis to a GT basis according to

$$\beta = \begin{pmatrix} \beta_{n-1,n-1} & \beta_{n-1,n-2} & \cdots & \beta_{n-1,0} \\ 0 & \beta_{n-2,n-2} & \cdots & \beta_{n-2,0} \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & \beta_{0,0} \end{pmatrix}, \quad (16)$$

where $\beta_{i,j} \in \{0,1\}$ for $i > j$, $\beta_{i,i} = 1$, and n is the number of qubits involved in the transformation. Following closely the notations used in [41], we define the following sets of indices for convenience, which we detail below. We note that all matrix operations are performed in modulo-2 space and the main diagonal elements are excluded when generating these sets.

- Update set $U(j)$: elements of this set are the row indices with non-zero entries in column j of the basis-transformation matrix β .
- Parity set $P(j)$: elements of this set are the column indices with non-zero entries in row j of the matrix $(\pi\beta^{-1} - \beta^{-1})$, where $\pi_{i,j} = 1$ if $i \leq j$, otherwise 0.
- Remainder set $R(j)$: elements of this set are the column indices with non-zero entries in row j of the matrix $\pi\beta^{-1}$.

The GT-based creation and annihilation operators are then

$$\begin{aligned} a_j^\dagger &\equiv \left[\sigma_x^{U(j)} \otimes \sigma_x^j \otimes \sigma_z^{P(j)} - i \sigma_x^{U(j)} \otimes \sigma_y^j \otimes \sigma_z^{R(j)} \right] / 2, \\ a_j &\equiv \left[\sigma_x^{U(j)} \otimes \sigma_x^j \otimes \sigma_z^{P(j)} + i \sigma_x^{U(j)} \otimes \sigma_y^j \otimes \sigma_z^{R(j)} \right] / 2, \end{aligned} \quad (17)$$

which can straightforwardly be shown to satisfy (2). We note that, e.g., the JW transformation is a special case of $\beta = \mathbf{1}$.

The state of the art implementation of the UCCSD ansatz [31] relies on the JW transformation and it considers a series of heuristics that were chosen carefully to optimize the resulting quantum circuits. To enable a proper comparison of those with our results then, it is necessary for us to consider a similar series of heuristics as well. Details of the heuristics we consider are provided in the SM Secs. B–G. Below, we briefly outline the steps in the order of applications. The cost function we consider for our concrete example is the number of two-qubit gates.

The outer-most loop of our approach considers different transformation matrices β . For a given mapping matrix β , we execute the following routines to construct and optimize our circuit that implements the UCCSD ansatz. The routines are designed to repeatedly call a suite of dedicated, automated circuit optimization tools, whose technical details may be found in [30, 31]. The efficiency of these tools allows us to quickly evaluate the cost function for different cases we consider in each of the subroutines.

Routine 1 *Fermionic level labeling*: Unlike in the JW transformation where $U(j)$ is an empty set and $P(j) = R(j)$, in the GT approach, myriads of combinations of sets $U(j)$, $P(j)$, and $R(j)$ are possible. To best take advantage of this, it is critical to carefully select which fermion level is mapped onto which qubit index. Exploring all possible mappings is however computationally prohibitively expensive. We thus resort to a simple greedy approach, whereby we explore one permutation at a time from a given fermion-level to qubit-index mapping. Specifically, from the given mapping, we apply the permutation that results in the most reduction in the quantum resource requirements. We iterate this process until no single permutation results in the reduction of quantum resource requirements. See SM Sec. B for detail. See **Subroutine 1** below for the cost function evaluation for each permutation.

Routine 2 *Inter-Trotter term ordering*: Demonstrated in [1, 14, 31, 38, 45, 46] was that ordering the Trotter terms appropriately can lead to large savings in the two-qubit gate counts due to gate cancellations between the neighboring Trotter term circuits. While a similar approach is indeed possible – and we base our method on the work reported in [31] – in our GT approach, a non-trivial modification to a method reported in [31] needs to be made. Specifically, we need to preprocess each Trotter term to determine its eligibility for being classified under the same equivalence class, the elements of which have the opportunities for resource savings when placed next to one another on a quantum circuit via the aforementioned gate cancellations. The eligibility criteria are straightforward for the JW transformation, considered in [31]. See SM Sec. D for details for the GT method. Once the equivalence classes according to the eligibility of each Trotter term is determined, we use a simple greedy approach to order the Trotter-term elements of each equivalence class to reduce the two-qubit gate counts.

Subroutine 1 *Intra-Trotter term ordering:* For a given fermion-label to qubit-index mapping, an efficient method to implement a single- or a double-fermion excitation Trotter term is known [47] in the case where the JW transformation is used, and an aggressively optimized circuit construction that leverages full qubit-to-qubit connectivity is available in [31]. The method in [31] relies on a careful ordering of the intra-Trotter operator implementation, where, for instance, an example double-fermion Trotter term $e^{t_{pqrs}(a_p^\dagger a_q^\dagger a_r a_s - h.c.)}$ is expanded to $\sigma_{x,y,z}$ -based intra-Trotter terms. To enable an efficient implementation in other transformations used in our GT approach, we compute the cost function for every possible permutation of the intra-Trotter terms (see SM Sec. C for detail). We choose an ordering with the least cost.

The resulting optimized circuit implements the UCCSD ansatz in the chosen transformation basis defined by β . However, with the exception of the JW transformation where $\beta = \mathbf{1}$, the GT β matrix requires us to implement the initial mapping of the basis at the beginning of the circuit. This incurs an overhead $O(n^2/\log(n))$ in the two-qubit gate counts [36]. To obtain the final quantum resource requirement, we call an automated optimizer with the input quantum circuit that consists of the prefix subcircuit that implements β and the postfix subcircuit that implements the β -basis UCCSD ansatz.

We note that in [31] the concept of “bosonic” excitations is discussed. Effectively, in the JW transformation, whenever a pair of neighboring qubits, whichever appropriate fermion levels they correspond to, are excited to yet another pair of neighboring qubits that denote another set of fermion levels, the circuit that implements such an excitation term can be dramatically simplified to require only two two-qubit gates, while requiring only half the number of qubits that would otherwise be required (see SM Sec. E for the cases where the pairs do not neighbor). To take advantage of this, we use a juxtaposition of the bosonic circuit written according to the JW transformation and non-bosonic circuit written according to our GT approach. We note in passing that, to return from the half-qubit space of the bosonic circuit to the full-qubit space of the non-bosonic circuit, at most $n/2$ CNOT gates are expended. All our circuit metrics appropriately reflect this.

Table 2 shows circuit metrics, measured according to the number of two-qubit gates used to implement the UCCSD ansatz circuit, for different molecules of our choice. We show the results for the JW, BK, and the best GT transformations that our heuristic toolchain specified above found for comparison. To find the best GT transformations, we used a particle swarm optimization, as detailed in SM Secs. F and G. The advantages offered by the GT transformations vary in the suite of molecules we consider, ranging from 1.44% to 21.43%. This demonstrates the capability of our heuristics that it is indeed possible to further optimize the quantum circuits over the previous state of the art obtained via the JW transformation by considering GT transformations, custom selected for different input cases.

We make an open-ended note here on the savings ratio obtained. The savings ratio for the GT method is obtained by comparison to the JW transformation, which often admits better circuit optimization over the BK transformation in the instances we considered, especially around the interesting regimes where the ansatz circuit is large enough to recover, e.g., chemical accuracy, but is sufficiently small to merit VQE implementations on a near-term, noisy device. In the asymptotic limit, the BK transformation, one of the GTs, is known to incur $O(\log(n))$ overhead, an exponential advantage over the JW transformation that incurs $O(n)$ overhead [48]. Provided that, in the small instances, the savings ratio can be larger than 20%, a question arises as to how the ratio will behave in the intermediate-

Molecule	n	NE	JW	BK	GT	Improve(%)
HF	6	3	30	29	25	16.67
LiH	6	3	30	29	25	16.67
BeH ₂	8	9	70	71	60	14.29
NH ₃	14	52	485	607	478	1.44
<hr/>						
H ₂ O(4)	8	4	42	50	33	21.43
H ₂ O(5)	8	5	44	52	35	20.45
H ₂ O(6)	8	6	46	47	37	19.56
H ₂ O(8)	10	8	68	88	63	7.35
H ₂ O(9)	10	9	71	89	66	7.04
H ₂ O(11)	10	11	93	110	87	6.45
H ₂ O(12)	10	12	95	112	89	6.32
H ₂ O(14)	10	14	114	140	111	2.63
H ₂ O(16)	10	16	135	166	131	2.96
H ₂ O(17)	10	17	137	168	133	2.91

Table 2: Number of two-qubit gates required for the VQE simulation of different molecules with different fermion to qubit transformations. n is the dimension of the β matrix. NE is the number of excitation terms considered in the UCCSD ansatz. JW/BK are the number of two-qubit gates with JW or BK transformations. GT is the number of two-qubit gates given by the best β other than JW or BK. All molecules use the STO-3G basis, and ansatz terms are determined according to the HMP2 ordering. (Top): Here, we report the results for the cases where the HMP2 method was used to reach chemical accuracy. (Bottom): Here, we report the results for the HMP2 progression for a water molecule. The number in the parentheses next to H₂O indicates the total number of excitation terms (D_α) considered for the UCCSD ansatz. The cutoff of H₂O(17) is chosen based on the recovery of chemical accuracy (see Fig. 4).

sized, aforementioned interesting regime. We mention here a couple possibilities. It may be that the improved circuit efficiency by the GT method for a particular excitation term could be offset by the increased cost of implementing other excitation terms. In this case, the ratio is expected to drop. It could also be that the JW transformation, due to its lack of nimbleness in the definition of transformation to help reduce potentially large cost of adding the excitation terms to the ansatz circuit, incurs a large overhead, thus resulting in the more significant savings ratio. Further research is required to provide better predictions. The low reduction rate observed for larger sized circuits in Table 2 may in part be due to the ineffectiveness of the classical optimization strategy and the limited classical computational resources used, which is further discussed in SM Sec. G.

5 Discussion

So far in this paper, in the optimization of the FT-regime circuits, we have considered efficient implementations of each Trotter term. It should however be noted that a parallel implementation of multiple Trotter terms should also be considered in the R_z gate depth reduction of the quantum simulation circuits. Based on our circuit construction detailed in Sec. 3, we propose the following methodology for optimization over the parallel implementation.

As discussed in Sec. 3, the circuit that implements the Trotter term (see Fig. 2) consists largely of three parts: an initial CNOT gate network that computes a linear function of

Boolean input variables, a triply-controlled R_x gate, and the inverse of the initial CNOT gate network. Denoting the Boolean variables at the input as, in the order from top to bottom, a , b , c , and d , the CNOT gate network outputs a , $b \oplus c$, $a \oplus c$, and $a \oplus d$. The bottom three outputs remain invariant over the action of the triply-controlled R_x gate. This means that any linear functions of $b \oplus c$, $a \oplus c$, and $a \oplus d$ are accessible for use in the implementation of CNOT gates that corresponds to the JW σ_z strings for other Trotter terms. This is so, because the invariance is required to implement the appropriate inverse of the CNOT gates that were used to take the JW σ_z strings into consideration in the first place, as per our circuit construction shown in SM Fig. 10. Heuristic methods that collect those Trotter terms and simultaneously implement them can then be used to optimize the depth of the quantum circuit.

For the pre-FT regime VQE simulations, we have proposed a general framework that leverages the predictive and corrective power of perturbation theory. The predictive feature of our framework resembles that of the ADAPT-VQE method [12], which aims to iteratively add terms to a generalized Trotterized exponential ansatz by following the steepest descent based on the calculations of numerical gradients. Despite trying to solve the same problem, our framework adopts a different philosophy and computes an approximate (perturbative) correction to the wavefunction and chooses the following operator based on how much of a correction an operator could potentially provide. Due to the two completely different approaches to the problem, the two methods could have vastly different use cases for which they are suitable respectively. To identify the difference in use cases, further thorough investigation into different areas of applications are needed which is beyond the scope of this paper.

Additionally, our framework is exceedingly flexible and provides means to balance between resource constraints and accuracy requirement. First of all, our framework not only works with the UCCSD ansatz, but any unitary ansatz, including the generalized Trotterized exponential ansatz used in the ADAPT-VQE method. Second, the predictive feature and the corrective feature can be used independently from each other. For instance, the perturbative corrections can be used along with the ADAPT-VQE method, which provides the ansatz term prediction. Note that doing so, according to our framework, does not incur additional cost in terms of circuit complexity; it incurs overhead only in terms of the number of measurements. Our specific examples described in this manuscript results in smaller number of measurements required, since both the prediction and the correction are obtained by the same set of measurements, as opposed to two different set of measurements that would otherwise be required, had we combined ADAPT-VQE with our perturbative correction. Third, the choice of the set of orbital sequence $\{\alpha'\}$ in (9) can be leveraged to balance between accuracy requirement and resource constraints. On one hand, the size of the set can be decreased to reduce the number of measurements needed. On the other hand, it is also possible to include orbital sequences outside of the set $\{\alpha\}$ for the ansatz operator in order to achieve better accuracy without sacrificing circuit complexity. For instance, one could include triple excitation sequences in $\{\alpha'\}$ along with the UCCSD ansatz to provide a perturbative ‘‘T’’ correction to the final ground state energy. In this case the increased accuracy does not correspond to increased circuit complexity but only incurs overhead in terms of the number of measurements.

We note that, extending and generalizing the framework used for considering the bosonic terms in the JW transformation and the non-bosonic terms in the GT transformation, the use of multiple fermion to qubit transformations β for a given set of excitation terms could be of value in reducing the overall resource requirement. Drawing from the fact that different β transformations result in different resource requirements in imple-

menting the excitation terms of a target UCCSD ansatz circuit, it is reasonable to expect that a certain subset of the excitation terms may be more efficiently implemented by one β transformation and some other excitation terms may be implemented more efficiently by yet another β transformation. Thus, dividing the set of excitation terms required for preparing a UCCSD ansatz state into subsets of excitation terms that may be more efficiently implemented by respective, appropriate choices of β transformations for each of the subsets may prove to be more advantageous in the quantum resource requirement. Optimizing over the tug of war expected between the overhead cost incurred due to the switching of the transformations and the savings obtained via the tailor-made choices of the transformations remains as future work.

6 Conclusion

Quantum simulations performed by quantum computers have long been thought to be one of the most promising quantum applications that will prove advantageous over classical computers. Despite the recent technological advancements made by the community, there still remains a gap between what a quantum device can realistically achieve and what is required to demonstrate a practical advantage in running a quantum simulation on a quantum device. In an attempt to bridge this gap, in this paper, we focused on optimizing the quantum resource requirements for both the FT and pre-FT regime quantum simulations of fermionic systems. Our FT-regime optimized approach yields a R_z -depth of one quantum circuit for each Trotterized evolution operator with only two R_z gates to implement in total, while requiring just a single ancilla qubit. Our pre-FT regime optimized approach leverages perturbation theory, capable of determining the best ansatz terms to consider to enable a rapid convergence to the sought-after ground-state energy of a simulated molecule that in turn results in the reduction of gate counts required to reach a certain accuracy threshold. We also considered GT methods to obtain significant quantum resource savings, upwards of 20% in some instances, over the conventional JW or BK transformations in practice. While these results were applied to a specific set of fermionic systems in this paper as a concrete example, we expect similarly-spirited works could leverage our methodologies. We believe the savings in the quantum resource requirement demonstrated by our approaches help bring the day of solving practical problems using a quantum computer closer.

Acknowledgement

This work was supported, in part, by the DOE BES QIS Program (de-sc0019449). The authors thank Diego Acosta and Germán Mauricio Muñoz at IonQ, Inc., for their efforts in the implementation of computing systems used to generate the data presented in this manuscript. The authors further thank Prof. Reinhold Blümel for helpful discussion.

References

- [1] Ryan Babbush, Jarrod McClean, Dave Wecker, Alán Aspuru-Guzik, and Nathan Wiebe. Chemical basis of trotter-suzuki errors in quantum chemistry simulation. *Physical Review A*, 91(2), Feb 2015. ISSN 1094-1622. DOI: [10.1103/physreva.91.022311](https://doi.org/10.1103/physreva.91.022311).
- [2] Rodney J. Bartlett, Stanislaw A. Kucharski, and Jozef Noga. Alternative coupled-

- cluster ansätze ii. the unitary coupled-cluster method. *Chemical Physics Letters*, 155 (1):133 – 140, 1989. ISSN 0009-2614. DOI: [10.1016/S0009-2614\(89\)87372-5](https://doi.org/10.1016/S0009-2614(89)87372-5).
- [3] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders. Efficient quantum algorithms for simulating sparse hamiltonians. *Communications in Mathematical Physics*, 270(2):359–371, 2007. ISSN 1432-0916. DOI: [10.1007/s00220-006-0150-x](https://doi.org/10.1007/s00220-006-0150-x).
- [4] Alex Bocharov, Martin Roetteler, and Krysta M. Svore. Efficient synthesis of universal repeat-until-success quantum circuits. *Physical Review Letters*, 114:080502, Feb 2015. DOI: [10.1103/PhysRevLett.114.080502](https://doi.org/10.1103/PhysRevLett.114.080502).
- [5] Sergey B. Bravyi and Alexei Yu. Kitaev. Fermionic quantum computation. *Annals of Physics*, 298(1):210 – 226, 2002. ISSN 0003-4916. DOI: [10.1006/aphy.2002.6254](https://doi.org/10.1006/aphy.2002.6254).
- [6] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38):9456–9461, 2018. ISSN 0027-8424. DOI: [10.1073/pnas.1801723115](https://doi.org/10.1073/pnas.1801723115).
- [7] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74:4091–4094, May 1995. DOI: [10.1103/PhysRevLett.74.4091](https://doi.org/10.1103/PhysRevLett.74.4091).
- [8] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski. Cloud quantum computing of an atomic nucleus. *Physical Review Letters*, 120:210501, May 2018. DOI: [10.1103/PhysRevLett.120.210501](https://doi.org/10.1103/PhysRevLett.120.210501).
- [9] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982. ISSN 1572-9575. DOI: [10.1007/BF02650179](https://doi.org/10.1007/BF02650179).
- [10] Craig Gidney. Halving the cost of quantum addition. *Quantum*, 2:74, Jun 2018. ISSN 2521-327X. DOI: [10.22331/q-2018-06-18-74](https://doi.org/10.22331/q-2018-06-18-74).
- [11] Pranav Gokhale, Olivia Angiuli, Yongshan Ding, Kaiwen Gui, Teague Tomesh, Martin Suchara, Margaret Martonosi, and Frederic T Chong. Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families. *arXiv preprint arXiv:1907.13623*, 2019.
- [12] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications*, 10(1):3007, Jul 2019. ISSN 2041-1723. DOI: [10.1038/s41467-019-10988-2](https://doi.org/10.1038/s41467-019-10988-2).
- [13] Emily Grumblin and Mark Horowitz, editors. *Quantum Computing: Progress and Prospects*. The National Academies Press, Washington, DC, 2019. ISBN 978-0-309-47969-1. DOI: [10.17226/25196](https://doi.org/10.17226/25196).
- [14] Matthew B. Hastings, Dave Wecker, Bela Bauer, and Matthias Troyer. Improving quantum algorithms for quantum chemistry. *Quantum Info. Comput.*, 15(1–2):1–21, January 2015. ISSN 1533-7146. DOI: [10.5555/2685188.2685189](https://doi.org/10.5555/2685188.2685189).
- [15] Cornelius Hempel, Christine Maier, Jonathan Romero, Jarrod McClean, Thomas Monz, Heng Shen, Petar Jurcevic, Ben P. Lanyon, Peter Love, Ryan Babbush, Alán Aspuru-Guzik, Rainer Blatt, and Christian F. Roos. Quantum chemistry calculations on a trapped-ion quantum simulator. *Physical Review X*, 8:031022, Jul 2018. DOI: [10.1103/PhysRevX.8.031022](https://doi.org/10.1103/PhysRevX.8.031022).
- [16] Mark R. Hoffmann and Jack Simons. A unitary multiconfigurational coupled-cluster method: Theory and applications. *The Journal of Chemical Physics*, 88(2):993–1002, 1988. DOI: [10.1063/1.454125](https://doi.org/10.1063/1.454125).
- [17] Zhang Jiang, Kevin J. Sung, Kostyantyn Kechedzhi, Vadim N. Smelyanskiy, and Sergio Boixo. Quantum algorithms to simulate many-body physics of correlated fermions.

- Physical Review Applied*, 9(4), Apr 2018. ISSN 2331-7019. DOI: [10.1103/physrevapplied.9.044036](https://doi.org/10.1103/physrevapplied.9.044036).
- [18] P. Jordan and E. Wigner. Über das paulische äquivalenzverbot. *Zeitschrift für Physik*, 47(9):631–651, 1928. ISSN 0044-3328. DOI: [10.1007/BF01331938](https://doi.org/10.1007/BF01331938).
- [19] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017. ISSN 1476-4687. DOI: [10.1038/nature23879](https://doi.org/10.1038/nature23879).
- [20] James Kennedy and Russell C Eberhart. A discrete binary version of the particle swarm algorithm. In *1997 IEEE International conference on systems, man, and cybernetics. Computational cybernetics and simulation*, volume 5, pages 4104–4108. IEEE, 1997. DOI: [10.1109/ICSMC.1997.637339](https://doi.org/10.1109/ICSMC.1997.637339).
- [21] Ian D. Kivlichan, Craig Gidney, Dominic W. Berry, Nathan Wiebe, Jarrod McClean, Wei Sun, Zhang Jiang, Nicholas Rubin, Austin Fowler, Alán Aspuru-Guzik, and et al. Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization. *Quantum*, 4:296, Jul 2020. ISSN 2521-327X. DOI: [10.22331/q-2020-07-16-296](https://doi.org/10.22331/q-2020-07-16-296).
- [22] Joonho Lee, William J. Huggins, Martin Head-Gordon, and K. Birgitta Whaley. Generalized unitary coupled cluster wave functions for quantum computation. *Journal of Chemical Theory and Computation*, 15(1):311–324, Jan 2019. ISSN 1549-9618. DOI: [10.1021/acs.jctc.8b01004](https://doi.org/10.1021/acs.jctc.8b01004).
- [23] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996. ISSN 0036-8075. DOI: [10.1126/science.273.5278.1073](https://doi.org/10.1126/science.273.5278.1073).
- [24] Guang Hao Low and Isaac L. Chuang. Optimal hamiltonian simulation by quantum signal processing. *Physical Review Letters*, 118:010501, Jan 2017. DOI: [10.1103/PhysRevLett.118.010501](https://doi.org/10.1103/PhysRevLett.118.010501).
- [25] Guang Hao Low and Isaac L. Chuang. Hamiltonian Simulation by Qubitization. *Quantum*, 3:163, July 2019. ISSN 2521-327X. DOI: [10.22331/q-2019-07-12-163](https://doi.org/10.22331/q-2019-07-12-163).
- [26] D. Maslov, Y. Nam, and J. Kim. An outlook for quantum computing [point of view]. *Proceedings of the IEEE*, 107(1):5–10, Jan 2019. ISSN 1558-2256. DOI: [10.1109/JPROC.2018.2884353](https://doi.org/10.1109/JPROC.2018.2884353).
- [27] Dmitri Maslov. Advantages of using relative-phase toffoli gates with an application to multiple control toffoli optimization. *Physical Review A*, 93:022311, Feb 2016. DOI: [10.1103/PhysRevA.93.022311](https://doi.org/10.1103/PhysRevA.93.022311).
- [28] Dmitri Maslov and Yunseong Nam. Use of global interactions in efficient quantum circuit constructions. *New Journal of Physics*, 20(3):033018, mar 2018. DOI: [10.1088/1367-2630/aaa398](https://doi.org/10.1088/1367-2630/aaa398).
- [29] Yunseong Nam and Dmitri Maslov. Low-cost quantum circuits for classically intractable instances of the hamiltonian dynamics simulation problem. *npj Quantum Information*, 5(1):44, 2019. ISSN 2056-6387. DOI: [10.1038/s41534-019-0152-0](https://doi.org/10.1038/s41534-019-0152-0).
- [30] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information*, 4(1):23, 2018. ISSN 2056-6387. DOI: [10.1038/s41534-018-0072-4](https://doi.org/10.1038/s41534-018-0072-4).
- [31] Yunseong Nam, Jwo-Sy Chen, Neal C Pseni, Kenneth Wright, Conor Delaney, Dmitri Maslov, Kenneth R Brown, Stewart Allen, Jason M Amini, Joel Apisdorf, Aleksey Beck, Kristin M. Blinov, Vandiver Chaplin, Mika Chmielewski, Coleman Collins, Shantanu Debnath, Andrew M. Ducore, Kai M. Hudek, Matthew Keesan, Sarah M. Kreikemeier, Jonathan Mizrahi, Phil Solomon, Mike Williams, Jaime David Wong-Campos, Christopher Monroe, and Jungsang Kim. Ground-state energy estimation of

- the water molecule on a trapped-ion quantum computer. *npj Quantum Information*, 6(1):1–6, 2020. DOI: [10.1038/s41534-020-0259-3](https://doi.org/10.1038/s41534-020-0259-3).
- [32] Yunseong Nam, Yuan Su, and Dmitri Maslov. Approximate quantum fourier transform with $\mathcal{O}(n \log(n))$ t gates. *npj Quantum Information*, 6(1), Mar 2020. ISSN 2056-6387. DOI: [10.1038/s41534-020-0257-5](https://doi.org/10.1038/s41534-020-0257-5).
- [33] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. DOI: [10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667).
- [34] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis. Scalable quantum simulation of molecular energies. *Physical Review X*, 6:031007, Jul 2016. DOI: [10.1103/PhysRevX.6.031007](https://doi.org/10.1103/PhysRevX.6.031007).
- [35] Robert M. Parrish, Lori A. Burns, Daniel G. A. Smith, Andrew C. Simmonett, A. Eugene DePrince, Edward G. Hohenstein, Uğur Bozkaya, Alexander Yu. Sokolov, Roberto Di Remigio, Ryan M. Richard, Jérôme F. Gonthier, Andrew M. James, Harley R. McAlexander, Ashutosh Kumar, Masaaki Saitow, Xiao Wang, Benjamin P. Pritchard, Prakash Verma, Henry F. Schaefer, Konrad Patkowski, Rollin A. King, Edward F. Valeev, Francesco A. Evangelista, Justin M. Turney, T. Daniel Crawford, and C. David Sherrill. Psi4 1.1: An open-source electronic structure program emphasizing automation, advanced libraries, and interoperability. *Journal of Chemical Theory and Computation*, 13(7):3185–3197, 2017. DOI: [10.1021/acs.jctc.7b00174](https://doi.org/10.1021/acs.jctc.7b00174).
- [36] Ketan N. Patel, Igor L. Markov, and John P. Hayes. Optimal synthesis of linear reversible circuits. *Quantum Info. Comput.*, 8(3):282–294, March 2008. ISSN 1533-7146. DOI: [10.5555/2011763.2011767](https://doi.org/10.5555/2011763.2011767).
- [37] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1):4213, 2014. ISSN 2041-1723. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [38] David Poulin, Matthew B. Hastings, Dave Wecker, Nathan Wiebe, Andrew C. Doherty, and Matthias Troyer. The trotter step size required for accurate quantum simulation of quantum chemistry. *Quantum Info. Comput.*, 15(5–6):361–384, April 2015. ISSN 1533-7146. DOI: [10.5555/2871401.2871402](https://doi.org/10.5555/2871401.2871402).
- [39] Markus Reiher, Nathan Wiebe, Krysta M. Svore, Dave Wecker, and Matthias Troyer. Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences*, 2017. ISSN 0027-8424. DOI: [10.1073/pnas.1619152114](https://doi.org/10.1073/pnas.1619152114).
- [40] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology*, 4(1):014008, oct 2018. DOI: [10.1088/2058-9565/aad3e4](https://doi.org/10.1088/2058-9565/aad3e4).
- [41] Jacob T. Seeley, Martin J. Richard, and Peter J. Love. The bravyi-kitaev transformation for quantum computation of electronic structure. *The Journal of Chemical Physics*, 137(22):224109, 2012. DOI: [10.1063/1.4768229](https://doi.org/10.1063/1.4768229).
- [42] O. Shehab, K. Landsman, Y. Nam, D. Zhu, N. M. Linke, M. Keesan, R. C. Pooser, and C. Monroe. Toward convergence of effective-field-theory simulations on digital quantum computers. *Physical Review A*, 100:062319, Dec 2019. DOI: [10.1103/PhysRevA.100.062319](https://doi.org/10.1103/PhysRevA.100.062319).
- [43] Mark Steudtner and Stephanie Wehner. Fermion-to-qubit mappings with varying

- resource requirements for quantum simulation. *New Journal of Physics*, 20(6):063010, jun 2018. DOI: [10.1088/1367-2630/aac54f](https://doi.org/10.1088/1367-2630/aac54f).
- [44] Masuo Suzuki. General theory of fractal path integrals with applications to many-body theories and statistical physics. *Journal of Mathematical Physics*, 32(2):400–407, 1991. DOI: [10.1063/1.529425](https://doi.org/10.1063/1.529425).
- [45] Andrew Tranter, Peter J. Love, Florian Mintert, Nathan Wiebe, and Peter V. Coveney. Ordering of trotterization: Impact on errors in quantum simulation of electronic structure. *Entropy*, 21(12):1218, Dec 2019. ISSN 1099-4300. DOI: [10.3390/e21121218](https://doi.org/10.3390/e21121218).
- [46] Dave Wecker, Bela Bauer, Bryan K. Clark, Matthew B. Hastings, and Matthias Troyer. Gate-count estimates for performing quantum chemistry on small quantum computers. *Physical Review A*, 90(2), Aug 2014. ISSN 1094-1622. DOI: [10.1103/physreva.90.022305](https://doi.org/10.1103/physreva.90.022305).
- [47] James D. Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. Simulation of electronic structure hamiltonians using quantum computers. *Molecular Physics*, 109(5):735–750, 2011. DOI: [10.1080/00268976.2011.552441](https://doi.org/10.1080/00268976.2011.552441).
- [48] Ciyou Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on mathematical software (TOMS)*, 23(4):550–560, 1997. DOI: [10.1145/279232.279236](https://doi.org/10.1145/279232.279236).

SUPPLEMENTARY MATERIAL

A Circuits for the fault-tolerant regime

In this section, we show in detail the circuit implementation of the Trotter terms, typically encountered in quantum chemistry simulations, optimized for the FT regime. First, we show how to implement a triply-controlled R_x gate, needed for our two-body term circuit shown in Fig. 2 of the main text. Then, we compare the resulting quantum resource requirements of our method against those of the method shown in [39], which can be optimized by combining with the work discussed in [10, 29], in terms of circuit size, width, and depth.

Fig. 6(a) shows an R_z - and T-gate optimized implementation of a triply-controlled R_x gate. We use the so-called relative phase Toffoli gate investigated in [27] instead of a conventional Toffoli gate to reduce the T gate count, excluding that required for the R_z gates, to 16. An additional advantage is that a relative phase Toffoli gate with three controls does not require an ancilla qubit, whereas a conventional triply-controlled Toffoli gate, decomposed into the standard gate set of CNOT, single-qubit Clifford, and T gates, requires at least one ancilla. For completeness, we show in Fig. 7 the explicit construction of the triply-controlled, relative phase Toffoli gate, imported verbatim from [27], although other methods, sometimes with a different gate set, have been reported elsewhere [7, 28].

Inspecting Fig. 6(a), we notice that we need to apply the two R_z gates in series. To reduce the R_z depth then, we can introduce an ancilla qubit prepared to $|0\rangle$. An R_z -depth one construction of the triply-controlled R_x gate is shown in Fig. 6(b). Should the same angle rotations happen in multiple qubits at the same time, a possibility if a molecule of interest that is being simulated has symmetries, a weight-sum trick reported and used in [10, 29] can be employed to exponentially reduce the number of R_z gates at a modest cost in the number of qubits and T gates. Two R_z gates to be applied in parallel however do not benefit from this trick.

Having fully described our circuit that implements a two-body term, we are now ready to compare it with the state of the art in [39]. Figure 6 of the appendix of [39] shows a circuit that implements a two-body term, albeit with different coefficients for each of the eight Pauli strings obtained by expanding σ_{\pm} with $\sigma_{x,y}$ in a two-body term. Circuit-width wise, four ancilla qubits are used. R_z depth is one.

Since a two-body term of our interest does not incur different coefficients, the circuit shown in Figure 6 of the appendix of [39] can be further optimized using the aforementioned weight-sum trick in [10, 29]. Using various formulas reported in [29], we arrive for our purposes at the following requirements: at most 32 T gates to compute the weight, where four full and three half adders are used, and seven additional ancilla qubits for each adder. The number of R_z gates to be applied is then reduced to four. In total, the optimized circuit requires 32 T gates, 4 R_z gates in one layer, and eleven ancilla qubits, four from Fig. 6 of the appendix of [39] and seven from the weight-sum trick reported in [10, 29]. This may be compared with our R_z -depth optimal implementation that requires 16 T gates, two R_z gates in one layer, and one ancilla qubit.

We next briefly describe a single-body term, reported elsewhere [17, 21] before, for completeness. Our work reported above may be considered as a generalization of it to a multi-body term. Note, while we do not explicitly show an N -body term circuit as we do not need it for our purposes, an optimal construction can straightforwardly be derived using the work reported herein.

For the single-body operator $e^{-i\theta/2(\sigma_+\sigma_- + \text{h.c.})}$, in analogy to the two-body case shown

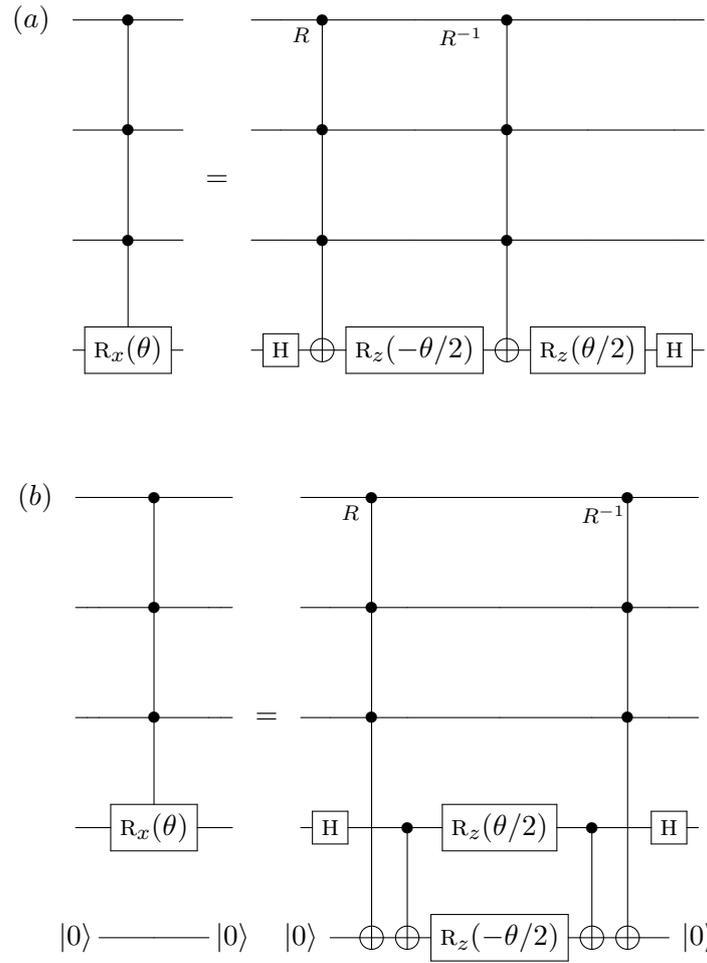


Figure 6: Triply-controlled $R_x(\theta)$ gate implementation with relative-phase triply-controlled NOT gates. The R -prepped triply-controlled Toffoli gate denotes a relative-phase Toffoli gate, i.e., the gate implements a triply-controlled-NOT gate up to relative phases. The R^{-1} -preppendix is used to denote its inverse. The method to implement the relative-phase Toffoli gate is detailed in [7, 27, 28].

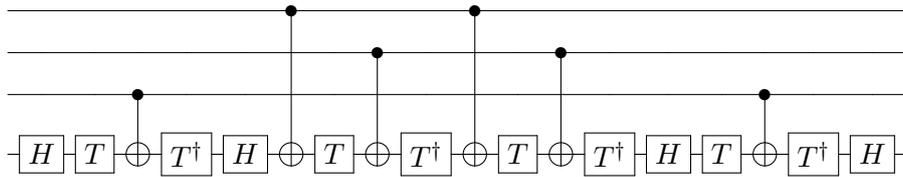


Figure 7: A quantum circuit for a relative-phase triply-controlled NOT gate. The circuit is taken verbatim from [27].

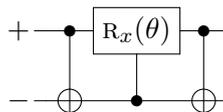


Figure 8: A circuit construction of a single-body term that is in analogy to the two-body circuit discussed earlier. The circuit is equivalent to that reported in Fig. 16 of [17] or Fig. 9 of [21] up to a basis transformation.

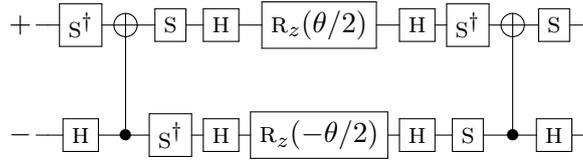


Figure 9: A depth-optimized construction of a single-body term in Fig. 8.

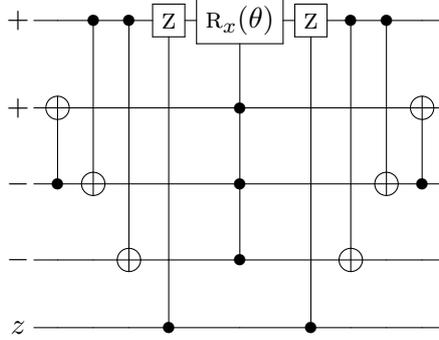


Figure 10: FT-regime optimized circuit for the two-body term $e^{-i\theta/2(\sigma_+\sigma_+\sigma_-\sigma_-\sigma_z+\text{h.c.})}$. We marked each qubit lines with +, -, or z to denote which qubits are associated with σ_+ , σ_- , or σ_z , respectively, for this particular example.

above, we may use the circuit shown in Fig. 8 to implement the operator. An R_z -depth optimized version is shown in Fig. 9. This results in a circuit that requires no ancilla qubit and two R_z gates in one layer.

We note in passing that if there are $\sigma_{x,y,z}$ operators that act on some other qubits, e.g., due to the JW transformation, we may simply modify the two-body circuit as follows. Consider, for example, there is a σ_z operator acting on an additional qubit. In this case, to implement $e^{-i\theta/2(\sigma_+\sigma_+\sigma_-\sigma_-\sigma_z+\text{h.c.})}$, we can modify the circuit in Fig. 2 and obtain the circuit shown in Fig. 10. A similar treatment is straightforwardly applicable for an arbitrary many body term.

B Fermionic level labeling

In this section, we detail the heuristic described in **Routine 1** in Sec. 4.2 which aims to reduce the number of two-qubit gates by exploiting different mappings from labels of fermion levels to qubit indices. Note that the Pauli strings for an excitation operator, after applying the fermion \mapsto qubit transformation of choice, depends on the qubit index values according to (17). A better optimized mapping from labels of fermion levels to qubit indices results in fewer Pauli matrices in a given Pauli string which leads to a smaller number of two-qubit gates. Notice that such a mapping is equivalent to maintaining a simple mapping from any fermion level k to qubit index k while rearranging the labeling of the fermion levels which is entirely arbitrary up until this point. Thus we may use the freedom of fermion level labeling to reduce the number of two-qubit gates required to implement the one- and two-body Trotter terms. For an n -qubit system representing n fermion levels, exploring all possible permutations of fermion level labeling is not computationally scalable as the total number of permutations is $n!$. We thus resort to a simple greedy approach as follows.

1. We first generate a set of unique permutation matrices $\{P_i\}$ for all possible k -swap operations. Here we define a k -swap operation as swapping k unique ordered pairs of indices $(2l_j, 2l_j + 1)$ with k other unique ordered pairs of indices $(2m_j, 2m_j + 1)$ with the same subscript index $j \in \{0, \dots, k - 1\}$ where $l_j \in \{0, \dots, n/2 - 2\}$ and $m_j \in \{l_j + 1, \dots, n/2 - 1\}$. We use $k = 2$ for all of the examples in the main text.
2. Denoting the initial fermion labels as a column vector L_0 , we can then go through all the permutation matrices in the first round and compute the number of two-qubit gates that corresponds to each instance of relabeled fermion levels $P_i L_0$ by simply adding the number of two-qubit gates required for each excitation term obtained from the intra-Trotter term reordering subroutine. When going through the permutation matrices, if any matrix results in a lower two-qubit gate count than all the previous matrices, we record that two-qubit gate count as N_1 and the corresponding relabeled level vector as L_1 .
3. We proceed with the greedy approach iteratively following the first round. Denoting the resulting relabeled fermion level vector from the m th round as L_m and the corresponding two-qubit gate count as N_m , we again go through all the permutation matrices and compute the corresponding two-qubit gate counts for each instance of relabeled levels $P_i L_m$. Taking $N_{m+1} = N_m$ initially, every time a matrix results in a two-qubit count lower than N_{m+1} , we record that gate count as N_{m+1} and the matrix as L_{m+1} . If there is not any matrix that results in a lower two-qubit count than N_m , we terminate the iteration and return L_m as the optimal labeling. Otherwise we proceed to the next round.

We note that once the labels are determined, they are applied to all relevant fermionic operators, including the molecular Hamiltonian, to be consistent throughout our simulation.

C Intra-Trotter term ordering

In this section, we provide the details of **Subroutine 1** described in Sec. 4.2 which aims to reduce the number of two-qubit gates by optimizing the ordering of the Pauli strings transformed from a two-body operator.

We start by briefly noting that each one-body term leads to only one ordering, up to inversion. This is so, because it contains only two Pauli strings. Therefore, the one-body term does not require a specific ordering.

We next consider two-body operators. A single two-body operator, after applying proper transformation and PF algorithm, contains eight subterms. Specifically, we have

$$U^{\text{two-body}} = \prod_{j=0}^7 e^{-i\theta \otimes_v \sigma^{(j,v)}/2}, \quad (18)$$

where j denotes the intra-Trotter term index, v denotes the qubit index, and $\sigma \in \{\mathbf{1}, \sigma_x, \sigma_y, \sigma_z\}$. Each of the intra-Trotter terms $e^{-i\theta \otimes_v \sigma^{(j,v)}/2}$ can be readily translated into a standard circuit as described in Eq. (8) in Sec. Methods H of [31]. In a naive implementation of U that uses an arbitrary fermion to qubit transformation, each intra-Trotter term results in $2(N_j - 1)$ number of CNOT gates, where N_j is the number of non-identity $\sigma^{(j,v)}$ in the j th Pauli string.

Compared to the JW transformation, in the GT method, it is useful to have a more concrete set of rules to determine which qubit may be used as a target qubit t , in relation

to the realization of U as a quantum circuit (see, e.g., Eq. 8 (b) in Sec. Methods H of [31], for the JW transformation). In the JW transformation, any choice of t such that $\sigma^{(j,t)}$ are either σ_x or σ_y , i.e., the qubits that correspond to the fermion labels in the two-body term, are good choices. However, if $\sigma^{(j,t)} = \mathbf{1}$ in any one of the eight Pauli strings in U above, which frequently occur in the GT method, the intended target qubit index t is unusable. Therefore, in the GT method, qubit choices t that lead to $\sigma^{(j,t)} = \mathbf{1}$ are removed from the set of eligible targets that start with all qubits that correspond to the fermion labels that appear in the two-body term of interest and will subsequently not be used in the inter-Trotter term reordering subroutine detailed in Sec. D.

Now, for a chosen target qubit of choice t , we consider the ordering of intra-Trotter terms to maximize the CNOT reduction. In particular, we only consider the CNOT reduction between adjacent Pauli string circuits. For each of the adjacent intra-Trotter terms, i.e., $e^{-i\frac{\theta}{2}\otimes_v\sigma^{(j,v)}}$ and $e^{-i\frac{\theta'}{2}\otimes_v\sigma^{(j+1,v)}}$, where $j \in [0, 6]$, we enumerate v through all non-target qubits. We then compare $\sigma^{(j,v)}$ and $\sigma^{(j+1,v)}$. For a particular control qubit v , if neither of the two aforementioned Pauli matrices $\sigma^{(j,v)}$ and $\sigma^{(j+1,v)}$ is $\mathbf{1}$, then the circuit can be expressed as Fig. 11. In this particular circuit, if we have $M_0 = M_2$, there is a two CNOT reduction, while, if we have $M_0 \neq M_2$, there is a one CNOT reduction, according to the circuit identity discussed in [31] Sec. Method H. Suppose now there are $m_j^{(t)}$ two-CNOT reductions and $n_j^{(t)}$ one-CNOT reductions for the chosen target qubit t . Then, the total number of CNOT reductions is $2m_j^{(t)} + n_j^{(t)}$ for the target qubit of choice t . The total optimized number of CNOTs for a chosen target t of a certain ordering is thus $N^{(t)} = \sum_{j=0}^7 2(N_j - 1) - \sum_{j=0}^6 (2m_j^{(t)} + n_j^{(t)})$.

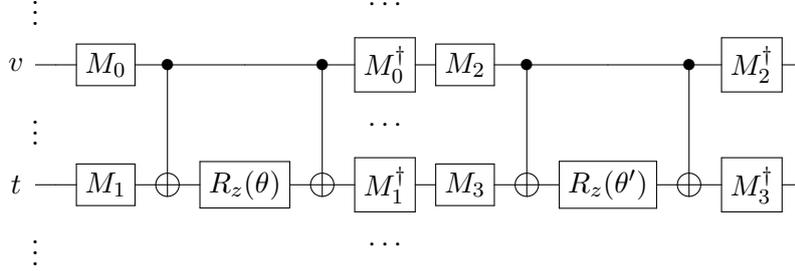


Figure 11: Example circuit for adjacent Pauli string circuits $e^{-i\frac{\theta}{2}\sigma_0^v\otimes\sigma_1^t}\dots e^{-i\frac{\theta'}{2}\sigma_2^v\otimes\sigma_3^t}\dots$. Subscripts l to Pauli operators are introduced to conveniently label the associated operators M_l shown in the figure. t denotes the target qubit, v denotes the control qubit, $\sigma_l \in \{\sigma_x, \sigma_y, \sigma_z\}$, and $M_l \in \{\mathbf{H}, \mathbf{s}^\dagger\mathbf{H}, \mathbf{1}\}$. If $\sigma_l = \sigma_x$, $M_l = \mathbf{H}$. If $\sigma_l = \sigma_y$, $M_l = \mathbf{s}^\dagger\mathbf{H}$. If $\sigma_l = \sigma_z$, $M_l = \mathbf{1}$.

At this point, we are equipped to compute and compare $N^{(t)}$ among all possible orderings and all possible targets to determine which one of the target qubit sets and the intra-Trotter orderings result in the minimum number of CNOTs. We note that for a certain ordering, there likely is a degeneracy in the choice of target qubit t that results in the same CNOT counts. All of these degenerate cases with the optimal resources are returned to the Inter-Trotter term ordering routine, detailed in Sec. D, for use.

D Inter-Trotter term ordering

In this subsection, we detail a heuristic to reduce the number of two-qubit gates by taking advantage of the freedom to arbitrarily order the Trotter terms. This is justified since any orderings respect the error bound provided by the PF algorithm. Denoted as **Routine**

2 in Sec. 4.2, this subroutine requires us to first run a preprocessing step based on the information passed from the intra-Trotter term optimization step described in Sec. C. Specifically, for operators $a_i^\dagger a_j$ and $a_i^\dagger a_j^\dagger a_k a_l$, we check if any one of the indices i and j for the former, and i, j, k , and l for the latter, may or may not be used as a target qubit in the circuit implementation of the operators in the standard compilation (see, e.g., Fig. 2d of [31] for a two-body term using the JW transformation). See Sec. C for details. We flag the qubit indices that cannot be used as a target as ineligible.

After all of the ineligibilities have been determined, we proceed according to a simple greedy approach. We first identify the most frequently eligible index, say p , across all terms. Then, we group all terms with the index p as an eligible target qubit and classify them under the equivalence class $[p]$. We next remove all the group elements from the list of one- and two-body operators. We repeat the procedure from the identification of the most frequent eligible index until no more operators are left in the list.

Note that the quantum resource cost reduction will likely now result in between the circuit representation of the elements of the same equivalence class, since the target qubit of two-qubit gates from each element of the same class is the same. Therefore, once all of the equivalence classes are specified, we consider permuting the orderings by which the elements are implemented on a quantum circuit. Considering all permutations can be prohibitively expensive. We thus use a simple greedy approach once more, first starting out with two of the elements that result in the most resource cost reduction. Then, we concatenate a next element, identified from the set of elements that have not been implemented in the circuit, based on the resource cost reduction. We repeat the concatenation process until no more element is left in the set. We note that each trial of testing out which element may be the best for the given iteration consists in general of four cases. This is so, since the circuit concatenation may be performed as a prefix or suffix, and the element to be concatenated can be considered in its original intra-term order or the reverse.

E Generalized Bosonic terms

The correspondence between Generalized Bosonic terms and Fermionic terms can be established as follows. A Fermionic double excitation term $\theta_{pqrs}(a_p^\dagger a_q^\dagger a_r a_s - h.c.)$, when transformed by the JW transformation, turns into $\theta_{pqrs}(\sigma_+^p \sigma_+^q \sigma_-^r \sigma_-^s \otimes_k \sigma_z^k - h.c.)$. If p and q belong to the same spatial orbital and r and s belong to yet another same spatial orbital, assuming no other terms that break the symmetry between p and q or r and s have been considered in the circuit, p and q levels may be encoded by a single qubit and likewise for the r and s levels. In this case, using p and r as representatives, the qubit-space operator can be simplified to $\theta_{pqrs}(\sigma_+^p \sigma_-^r \otimes_{k'} \sigma_z^{k'} - h.c.)$, where k' runs over the set of qubits σ_z operator needs to be applied for the given excitation term in the appropriately reduced space. The appropriately reduced space may include the single qubits that each denotes the reduced, symmetric levels and those that are not reduced. To illustrate, if two of the levels k_1 and k_2 in the original space require σ_z and if k_1 and k_2 are encoded into a single index k' , we simply call $\sigma_z^{k'}$ twice, one each for k_1 and k_2 . This for instance amounts to identity, which results in resource savings.

F Binary particle swarm optimization

In this section we detail the procedure for using binary PSO [20] to optimize the binary matrix β for GT. To encode the problem, we map the upper triangular entries in an $n \times n$ binary matrix β to a one dimensional binary vector $X \in \{0, 1\}^d$ where size $d = n(n-1)/2$.

Vector X then serves as the location vector for PSO. The cost function for PSO at a location X mapped from a β matrix is then defined as the number of two-qubit gates needed to implement a UCCSD evolution operator using GT with the β matrix and subsequently a series of heuristics described in Sec. 4.2.

We start by creating a swarm of particles, each with an initial location $X^i(t = 0)$ mapped from a β matrix sampled from the set of upper triangular matrices whose diagonal elements are ones and off-diagonal elements are zeros except for k of them. The non-negative integer t represents the time step and $t = 0$ is the beginning of the optimization. The total number of particles in the swarm is $\sum_{j=1}^k \binom{d}{j}$ where $\binom{d}{j}$ denotes a binomial coefficient. Specifically for our examples, we vary k from 1 to k_{\max} and we have $k_{\max} = 6$ for $n \leq 8$ and $k_{\max} = 3$ for the rest of the cases. Each particle also has a velocity vector $V^i(t)$ of the same dimension as $X^i(t)$. The initial matrix elements of the velocity vectors are taken to be zero.

At every time step t , we also keep track of a local optimal location vector $L^i(t)$ and the corresponding local optimal cost $s^i(t)$ for each particle i and a global optimal location vector $G(t)$. The local optimal location vector $L^i(t)$ is defined as the location vector with the lowest cost function value $s^i(t)$ across all the locations particle i has traveled to up to the time step t . The global location vector is the $L^i(t)$ with the lowest $s^i(t)$ among all the particles. The velocity vectors are updated from the t th step to the $(t + 1)$ th step according to

$$V_j^i(t + 1) = wV_j^i(t) + c_1(L_j^i(t) - X_j^i(t)) + c_2(G_j(t) - X_j^i(t)), \quad (19)$$

where w is the inertia parameter, c_1 is the cognitive parameter, and c_2 is the social parameter. For the examples here, we use various parameters with $w \in [-4, 4]$, $c_1 \in [0, 2]$, and $c_2 \in [0, 2]$. The subscript j denotes the j th element of the corresponding vector. The elements of the location vectors $X_j^i(t+1)$ are then updated to 0 if $\text{rand}() \leq 1/[1 + e^{-V_j^i(t+1)}]$ or 1 otherwise. Here $\text{rand}()$ is a real quasirandom number generating function that samples from the uniform distribution on $[0.0, 1.0]$.

In practice, binary PSO is usually terminated when t reaches a predetermined maximal time t_{\max} . A particular particle is also stopped when its location vector oscillates between two vectors for more than δt_{osc} times. To further save computation time spent on binary PSO, we impose a third rule that we stop a particle if its location vector maintains a Hamming distance larger than s from its initial location without reaching a lower two-qubit gate counts for more than δt_s steps. Such a termination condition allows similar search space sizes for β matrices of different sizes and optimization runs with different initial conditions. For our examples, we use $\delta t_{\text{osc}} = 10$, $s = 6$, $\delta t_s = 10$, $t_{\max} = 10000$ for $n \leq 8$, and $t_{\max} = 100$ for the rest of the cases. In most runs in our examples, our choice of t_{\max} is large enough that the optimization starts to oscillate before t reaches t_{\max} .

G General transformation results using binary particle swarm optimization

In this section we present and discuss examples of GT results using binary PSO with various amounts of classical computing resources. Specifically, we are interested in the fractional improvement in two-qubit gate count, ρ , defined according to

$$\rho = \frac{f_{\text{GT}}(R)}{f_{\text{GT}}(R) - f_{\text{JW}}}, \quad (20)$$

where f_{GT} and f_{JW} are the numbers of two-qubit gates needed to implement a UCCSD evolution operator using GT with binary PSO or JW, respectively. f_{GT} depends on the classical computing resources used by binary PSO which are determined mostly by the number of particles used since s and δt_s are the same for all optimization runs and t_{max} is always set to be sufficiently large for our examples as discussed in SM Sec. F. To contrast the amount of classical computing resources used with what is maximally possibly needed, we define a fractional classical computing resource metric as the ratio between the number of particles used, N , and the total number of possible variations of the β matrices, which can be written as

$$R = \frac{N}{2^{n(n-1)/2}}, \quad (21)$$

where n is the number of diagonal elements of the β matrices. Since the number of particles increase with k as discussed in SM Sec. F, we change k in our examples to effectively change the amount of classical computing resources used for binary PSO.

Figure 12 shows examples of the fractional improvement, ρ , as a function of the fractional classical computing resources used, R , for a few different molecules as well as for different numbers of ansatz terms used in the simulation of the water molecule. All the cases are simulated with the STO-3G basis set. We observe that ρ increases monotonically with R in all cases simulated. In some cases, *i.e.* hydrogen fluoride (HF), beryllium hydride (BeH₂), and water molecules with four to six ansatz terms included where n is small enough that we are able to explore a relatively large portion of possible variations of β matrices ($R > 10^{-3}$), we are able to obtain significant fractional improvement ranging from $\approx 14\%$ to over 20% . In other cases where n is large and we are only able to sample exponentially small portions of the location vector space in PSO, the fractional improvement obtained is much smaller and on the single digit percentage level. Thus we expect, based on the set of numerical results, that more cleverly exploring a large search space is likely a key to increasing the value of ρ ; Note the cost function $f_{\text{GT}}(\cdot)$ landscape is likely complex, making it tricky for binary PSO to avoid being trapped in local minima – We indeed observed many PSO jobs undergoing oscillatory behaviors.

H Qubit space reduction

In this section, we detail the procedures for the qubit space reduction (QSR) technique that optimizes the number of Pauli strings to be measured used in the main text. There are roughly two categories we consider: One is to treat the qubits that are in classically accessible states classically and the other is to take advantage of Bosonic states, *i.e.*, we expend only a single qubit for the two spin orbitals in the same spatial orbital with degenerate energy due to a particular choice of the ansatz.

As to the former, consider a Pauli string $S = A_0 A_1 \cdots A_{n-1}$ to be measured, where $A_i \in \{\sigma_x, \sigma_y, \sigma_z, I\}$, where I is a two-by-two identity operator. Conjugating this with a quantum state $|\Psi\rangle = |\psi\rangle|\phi\rangle$, where, without loss of generality, $|\psi\rangle$ denotes a quantum state of the qubits that are entangled due to the ansatz and $|\phi\rangle$ denotes a quantum state of the qubits that are in a classically accessible state, such as a computational basis state, we may write

$$\begin{aligned} \langle \Psi | S | \Psi \rangle &= \langle \psi | \langle \phi | \bigotimes_{i \in P} A_i \bigotimes_{j \in Q} A_j | \psi \rangle | \phi \rangle \\ &= \langle \psi | \bigotimes_{i \in P} A_i | \psi \rangle \langle \phi | \bigotimes_{j \in Q} A_j | \phi \rangle, \end{aligned} \quad (22)$$

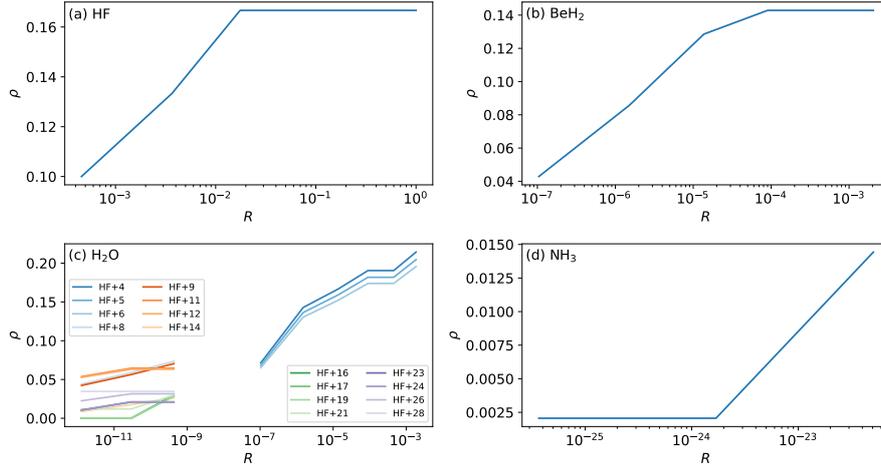


Figure 12: Fractional improvement ρ as a function of R , the fractional classical computing resources used for binary PSO. (a) HF, (b) BeH₂, (c) H₂O, and (d) NH₃. All molecules are simulated with the STO-3G basis set using the HMP2 method with UCCSD ansatz. For HF, BeH₂, and NH₃, the number of excitation terms included in the final UCCSD ansatz operators are 3, 9, and 52, respectively, which result in their respective ground state energy estimates within chemical accuracy from their known ground state energies. The size of the β matrices given by their respective number of diagonal terms n are shown in Table 2.

where we used sets P and Q to denote the set of qubits that are entangled and the set of qubits that are in a classically accessible state, respectively. Clearly, the second tensor product in the last line of (22) can be classically efficiently computed. Thus, we may evaluate just the first tensor product on a quantum computer and this amounts to a reduction in the number of Pauli strings to measure since there often are multiple second tensor products with the same first tensor product in the full set of Pauli strings to measure.

As to the latter, we start by a simple example to aid the description. Consider a quantum state $|\Psi\rangle$ defined according to

$$|\Psi\rangle = \sum_m |\psi_m\rangle (c_{m,0}|00\rangle + c_{m,1}|11\rangle). \quad (23)$$

Since the separated-out, two-qubit states $|00\rangle$ and $|11\rangle$ do not encode any more than a single qubit of information, this may be compressed to

$$|\Psi_{\text{comp}}\rangle = \sum_m |\psi_m\rangle (c_{m,0}|0\rangle + c_{m,1}|1\rangle). \quad (24)$$

Consider now a Pauli string, living in the full space spanned by $|\Psi\rangle$ in (23), $S = A_0 A_1 \cdots A_{n-1}$ to measure. Without loss of generality we designate $A_0 A_1$ as the Pauli product that lives in the separated-out two-qubit state space in (23). Conjugating S with $|\Psi\rangle$, we obtain

$$\begin{aligned} \langle \Psi | S | \Psi \rangle &= \sum_{m,n} \langle \psi_m | \bigotimes_{i>1} A_i | \psi_n \rangle (c_{m,0}^* c_{n,0} \langle 00 | A_0 A_1 | 00 \rangle + c_{m,0}^* c_{n,1} \langle 00 | A_0 A_1 | 11 \rangle \\ &\quad + c_{m,1}^* c_{n,0} \langle 11 | A_0 A_1 | 00 \rangle + c_{m,1}^* c_{n,1} \langle 11 | A_0 A_1 | 11 \rangle). \end{aligned} \quad (25)$$

Consider further a compressed Pauli string $S_{\text{comp}} = A_{\text{comp}} A_2 \cdots A_{n-1}$, living in the compressed space spanned by $|\Psi_{\text{comp}}\rangle$. Conjugating S_{comp} with $|\Psi_{\text{comp}}\rangle$, we obtain

$$\begin{aligned} \langle \Psi | S_{\text{comp}} | \Psi \rangle &= \sum_{m,n} \langle \psi_m | \bigotimes_{i>1} A_i | \psi_n \rangle (c_{m,0}^* c_{n,0} \langle 0 | A_{\text{comp}} | 0 \rangle + c_{m,0}^* c_{n,1} \langle 0 | A_{\text{comp}} | 1 \rangle \\ &\quad + c_{m,1}^* c_{n,0} \langle 1 | A_{\text{comp}} | 0 \rangle + c_{m,1}^* c_{n,1} \langle 1 | A_{\text{comp}} | 1 \rangle). \end{aligned} \quad (26)$$

A_0	A_1	A_{comp}
I	I	I
I	σ_x	0
I	σ_y	0
I	σ_z	σ_z
σ_x	I	0
σ_x	σ_x	σ_x
σ_x	σ_y	σ_y
σ_x	σ_z	0
σ_y	I	0
σ_y	σ_x	σ_y
σ_y	σ_y	$-\sigma_x$
σ_y	σ_z	0
σ_z	I	σ_z
σ_z	σ_x	0
σ_z	σ_y	0
σ_z	σ_z	I

Table 3: Conversion table from A_0A_1 to A_{comp} useful for reducing the number of measurements required for VQE simulations. 0 denotes a null matrix.

Inspecting (25) and (26), we observe that all that we need to satisfy for the two expressions to agree are

$$\langle 00|A_0A_1|00\rangle = \langle 0|A_{\text{comp}}|0\rangle, \quad (27)$$

$$\langle 00|A_0A_1|11\rangle = \langle 0|A_{\text{comp}}|1\rangle, \quad (28)$$

$$\langle 11|A_0A_1|00\rangle = \langle 1|A_{\text{comp}}|0\rangle, \quad (29)$$

$$\langle 11|A_0A_1|11\rangle = \langle 1|A_{\text{comp}}|1\rangle. \quad (30)$$

For all possible A_0 and A_1 , we report in Table 3 the conversion table for A_{comp} . Numerous null matrices that appear in Table 3, together with the reduced set of operators to measure that consists of three single-qubit Pauli matrices and an identity matrix, significantly reduces the measurement overhead.

I Number of two-qubit gate distribution

In this section we briefly discuss the distribution of extra CNOTs required when using GC with QSR(see Sec. 4.1.2). Note in GC one creates groups of multiple Pauli strings, drawn from the original set of strings required by a VQE simulation, that can simultaneously be measured. Each group can require a different number of extra CNOTs. It also requires a different number of measurements, typically smaller than the number of Pauli strings inside the group. To illustrate these requirements, in Fig. 13 we show the distribution of the number of extra CNOTs in the number of required measurements for the GC method for the water molecule example we use in the main text. We observe that, as expected, the number of extra CNOTs and the number of measurements to perform generally increases as a function of a larger ansatz size.

In the main text Fig. 5, we opt to use the average number of extra CNOTs as opposed to the maximal number of extra CNOTs for a given ansatz. This is so, since we expect the quantum computational error expected from each application of a two-qubit gate to even

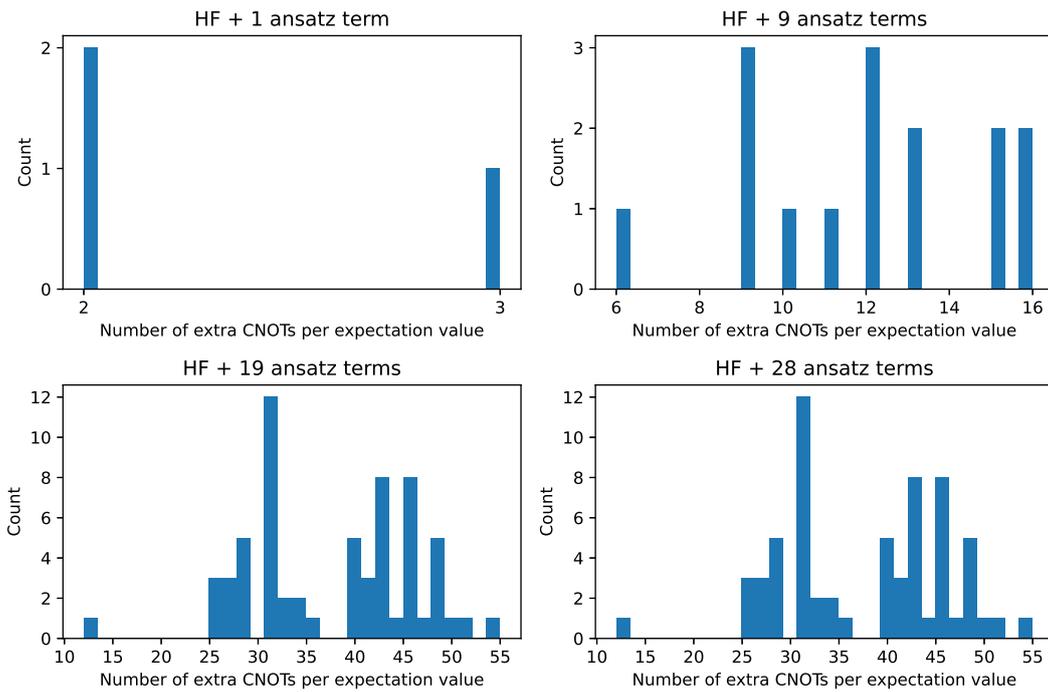


Figure 13: Histogram of the extra number of CNOTs required for the GC method per measurement value. The Pauli strings to be measured come from the Hamiltonian in HMP2 in (12). Shown are the cases for HF+1, HF+9, HF+19 and HF+28 for the water molecule example used in the main text.

out between the circuits that require a smaller than average number of extra CNOTs and the circuits that require a larger than average number of extra CNOTs. Thus we select the average number as a representative number that may best reflect a practical setting.