

ARTICLE OPEN

Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm

Bryan T. Gard^{1*}, Linghua Zhu¹, George S. Barron¹, Nicholas J. Mayhall², Sophia E. Economou¹ and Edwin Barnes¹

The variational quantum eigensolver is one of the most promising approaches for performing chemistry simulations using noisy intermediate-scale quantum (NISQ) processors. The efficiency of this algorithm depends crucially on the ability to prepare multi-qubit trial states on the quantum processor that either include, or at least closely approximate, the actual energy eigenstates of the problem being simulated while avoiding states that have little overlap with them. Symmetries play a central role in determining the best trial states. Here, we present efficient state preparation circuits that respect particle number, total spin, spin projection, and time-reversal symmetries. These circuits contain the minimal number of variational parameters needed to fully span the appropriate symmetry subspace dictated by the chemistry problem while avoiding all irrelevant sectors of Hilbert space. We show how to construct these circuits for arbitrary numbers of orbitals, electrons, and spin quantum numbers, and we provide explicit decompositions and gate counts in terms of standard gate sets in each case. We test our circuits in quantum simulations of the H_2 and LiH molecules and find that they outperform standard state preparation methods in terms of both accuracy and circuit depth.

npj Quantum Information (2020)6:10; <https://doi.org/10.1038/s41534-019-0240-1>

INTRODUCTION

Quantum simulation of fermionic systems, such as molecules, is one of the first envisioned applications of quantum computers, as famously proposed by Feynman.¹ The first protocol introduced for quantum chemistry simulations is based on the so-called phase estimation algorithm.^{2,3} This algorithm however requires a large number of quantum gates, leading to long quantum circuits that are challenging for existing and near-term noisy intermediate scale quantum (NISQ) devices.^{4,5} For such devices, alternative, hybrid algorithms are instead envisioned. In such algorithms, the work is shared between a quantum processor and a classical computer. In particular, the variational quantum eigensolver (VQE), first introduced and demonstrated experimentally by Peruzzo et al.,⁶ has become the prevailing algorithm for chemistry simulations with NISQ devices, with several milestone papers demonstrating the calculation of molecular energies and wavefunctions.^{7–12}

The VQE algorithm relies on preparing and measuring multi-qubit states based on a variational ansatz, and using the classical computer to optimize and update the variational parameters in this ansatz. Some of the advantages of VQE are that its variational character can provide some degree of error mitigation in the gates,^{7,13–15} and that it features shallower circuits compared to the phase estimation algorithm. The form of the ansatz is a crucial ingredient of VQE and one that can determine its success on NISQ devices. There are two main approaches in determining the ansatz. One approach is based on a technique from chemistry, the unitary coupled cluster method,^{16–18} translated into quantum gates by Trotterization.¹⁸ This approach tends to lead to deeper circuits than what is currently feasible on hardware, and is generally not exact. To address these issues, a new, iterative algorithm termed ADAPT-VQE, was recently put forward and was shown to enable a much more compact ansatz while simultaneously exhibiting higher accuracy.¹⁹ An alternative approach is to base the ansatz on the capabilities of the hardware and prepare states by combining parameterized gates available on the

processor.^{8,20} Such an ansatz has the advantage of compatibility with the capabilities of the hardware, and as such is NISQ-friendly. On the other hand, in its simplest form it is an ad hoc ansatz that can cause the algorithm to get stuck on ‘barren plateaus’²¹ as the number of qubits increases and the Hilbert space correspondingly grows exponentially. Therefore, for hardware-based ansätze to be a viable approach for problems of interest, they must be selected in a way that guarantees they span the part of the Hilbert space where the solution lives, while avoiding generating unphysical states.

Two ways to guarantee that the desired part of the Hilbert space is accessed include adding terms in the VQE energy function that penalize symmetry violations^{22,23} or carefully designing state preparation circuits so that they only produce states with the appropriate symmetries regardless of how their variational parameters are chosen. An early step toward the latter direction was taken by Wang et al.,²⁴ who focused on the preparation of states with a well-defined number of occupied spin orbitals and showed that the number of CNOT gates required for this scales polynomially with the number of qubits in the limit where the number of electrons is much smaller than the number of qubits. More recently, Barkoutsos et al.¹⁵ enforced particle number conservation by using the particle-hole representation in conjunction with a parametrized particle-conserving exchange-type gate,^{25–28} which we also make use of here. However, important open questions remain, including how other symmetries can also be built into the circuits and whether more efficient circuits containing the minimal number of parameters necessary to span the symmetry subspace exist.

In this paper, we address these questions by introducing state preparation circuits that provide a systematic, economical way to generate states with well-defined symmetries, including particle number, total spin, spin projection, and time-reversal. Our circuits incorporate the minimal number of parameters needed to fully span the appropriate symmetry subspace while avoiding all states outside this subspace. This general approach has two key

¹Department of Physics, Virginia Tech, Blacksburg, VA 24061, USA. ²Department of Chemistry, Virginia Tech, Blacksburg, VA 24061, USA. *email: bgard1@vt.edu

advantages: the first is that the true ground state is guaranteed to be contained within the space of states spanned by the circuit, and the second is that resources are not spent on generating irrelevant parts of the Hilbert space, reducing the complexity of the classical optimization step of the VQE algorithm. Eliminating extraneous parameters can dramatically speed up the optimization process and suppress the probability of getting trapped in local extrema or barren plateaus. We present circuits for arbitrary numbers of single-particle orbitals and electrons and for arbitrary spin quantum numbers. Our circuits are constructed with hardware constraint considerations, including a reduced number of CNOT gates that need only be applied between adjacent qubits in a linear array, making our work particularly suited for NISQ devices. In addition, since our construction conserves number and spin symmetries, symmetry verification techniques can be used to mitigate any errors which violate these symmetries.^{13,14,27} Our most general circuits, which conserve particle number, time-reversal, and spin symmetries, are guaranteed to span the appropriate symmetry subspace by construction. We also present circuits that conserve particle number and time-reversal symmetries, but not total spin; the symmetry-preserving properties of these circuits are verified with numerical calculations. We test the performance of our circuits against standard state preparation ansätze by running VQE simulations of the H_2 and LiH molecules. We find that our circuits outperform the standard methods in terms of both accuracy and circuit depth.

The paper is organized as follows. In the Results section, we define the basic gate set we use to manipulate particle number on the quantum processor and show how these gates can be systematically assembled into circuits that preserve particle number symmetry. We also discuss how to respect time-reversal symmetry as well. In addition, in the Results section, we show how to create circuits that respect total spin and spin projection symmetries in addition to particle number. We give some concluding remarks in the Discussion section. An appendix contains additional details about gate decompositions.

RESULTS

Particle number and time-reversal symmetries

In this work, we focus on mapping chemistry problems onto quantum processors using the Jordan-Wigner mapping,^{3,29} in which each qubit in the quantum processor corresponds to a particular spin-orbital, and the qubit states $|0\rangle$ and $|1\rangle$ encode the occupation of that spin-orbital. Any multi-electron state involving n spin-orbitals on the chemistry side can be mapped to a corresponding state of n qubits on the quantum processor. In this mapping, fixing the total number of electrons is tantamount to fixing the total number of qubits that are in the excited state $|1\rangle$. Thus, the Jordan-Wigner mapping relates fixed-particle-number subspaces to fixed-excitation subspaces in the qubit Hilbert space.

Formally, we can define the qubit subspace corresponding to m electrons occupying n spin-orbitals as

$$H_{n,m} = \text{span} \left\{ |s_1, s_2, \dots, s_n\rangle \mid s_i \in \{0, 1\}, \sum_{i=1}^n s_i = m \right\}. \quad (1)$$

This is the subspace of multi-qubit states containing m qubits in the $|1\rangle$ state and $n - m$ qubits in the $|0\rangle$ state, so that $\dim(H_{n,m}) = \binom{n}{m}$. A general state in this subspace can be represented as an arbitrary superposition of these basis states with complex coefficients and is thus characterized by $2 \dim(H_{n,m}) - 2$ real parameters, where we have removed two parameters by fixing the normalization and neglecting a global phase. In the absence of any other symmetries, $2 \dim(H_{n,m}) - 2$ is the minimal number of real variational parameters needed to prepare arbitrary trial states describing the correct number of electrons.

An additional symmetry that often arises in chemistry problems is time-reversal symmetry. This symmetry is typically present, for example, when one wants to solve the stationary Schrödinger equation in the absence of any applied magnetic field. In this case, one can always choose the energy eigenstates to be strictly real functions. Under the Jordan-Wigner mapping, this means that the coefficients appearing in the multi-qubit superposition states we prepare as trial states for the VQE algorithm should be restricted to real values. This will reduce the dimensionality of the target symmetry subspace by a factor of two down to $\dim(H_{n,m}) - 1$. Imposing this restriction on the trial states will prevent the classical optimizer from wasting time exploring a large portion of Hilbert space that does not contain any of the desired energy eigenstates. When we introduce our state preparation circuits below, we will see that time-reversal symmetry can be imposed easily after other symmetries are already built into the circuits, essentially just by fixing half of the variational parameters in the circuits in such a way that the resulting states are strictly real.

Before we introduce our general scheme for constructing particle-number-conserving state preparation circuits, we first present a few simple examples that may provide some intuition about the general structure of such circuits. First note that the cases $m = 0$ and $m = n$ are trivial since in each case, there is only a single state spanning the subspace. Therefore, we restrict our attention to $0 < m < n$ throughout this work. The simplest nontrivial example is the case of $m = 1$ electron in $n = 2$ orbitals. (Since we are presently only concerned with particle-number symmetry, these could be spin-orbitals or spatial orbitals in the case of spinless fermions. Spin symmetries will be incorporated in the next section). A circuit that spans the corresponding subspace $H_{2,1}$ is shown in Fig. 1.

This circuit requires only two parameters to span the single-excitation (one-electron) subspace, which is comprised of states of the form $\alpha|01\rangle + \beta|10\rangle$. Here, although α and β are complex, they contribute only two real parameters after we impose normalization $|\alpha|^2 + |\beta|^2 = 1$ and discard a global phase. The key ingredient in this circuit is a two-qubit entangling gate that we have denoted as $A(\theta, \phi)$. In the basis $|00\rangle, |01\rangle, |10\rangle, |11\rangle$, it is defined as¹⁵

$$A(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & e^{i\phi} \sin \theta & 0 \\ 0 & e^{-i\phi} \sin \theta & -\cos \theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2)$$

It is clear from the form of this exchange-type gate that it preserves particle number since it mixes $|01\rangle$ and $|10\rangle$ but does nothing to the $|00\rangle, |11\rangle$ subspace. The initial X gate on the first qubit in Fig. 1 brings the two-qubit state into the one-excitation subspace, while the subsequent A gate generates all possible superpositions within this subspace, as can be seen by inspection in this case. If we wish to also impose time-reversal symmetry, then it suffices to set $\phi = 0$ in each of the A gates. This removes the phase from each coefficient in the resulting superposition state without restricting the magnitude, thus ensuring that the resulting state is an arbitrary real state. The A gate plays a central role in our state preparation circuits. It can be decomposed into a sequence of two single-qubit gates and three CNOT gates, as




Fig. 1 2-Qubit A circuit. A simple example of a 2 qubit circuit which exactly spans the subspace defined by 1 excitation, $\alpha|01\rangle + \beta|10\rangle$, with two parameters (θ, ϕ) . This circuit saturates the lower bound on the number of real parameters required to construct arbitrary states in $H_{2,1}$.




Fig. 2 A gate decomposition. Decomposition of the A gate in terms of elementary single and two-qubit gates. $R(\theta, \phi) = R_z(\phi + \pi)R_y(\theta + \pi/2)$, where $R_z(\theta) = \exp(-i\theta\sigma_z/2)$, $R_y(\phi) = \exp(-i\phi\sigma_y/2)$.




Fig. 3 4-Qubit A circuit. An example circuit for the case of $n = 4$, $m = 2$ which exactly spans the subspace defined by six basis states using the minimal number (10) of parameters. Note that ϕ_1 is used three times.

shown in Fig. 2. This decomposition is minimal in the number of CNOT gates. In operator form, the A gate is

$$A = e^{-\frac{i}{2}(\theta_2 - \phi)Z_2} e^{-\frac{i}{2}(\theta X_1 X_2 + \theta Y_1 Y_2 + \frac{1}{2}Z_1 Z_2)} e^{-\frac{i}{2}(\theta Z_1 - \phi Z_2)},$$

up to an irrelevant global phase.³⁰

Building upon the example of Fig. 1, we find that we can generate trial states corresponding to other particle and orbital numbers using a similar construction.

Figure 3 shows another example circuit for the case of $m = 2$ electrons in $n = 4$ orbitals. In this case, we start with two X gates to bring the state into the two-excitation subspace, and we then apply a series of A gates between neighboring qubits to create arbitrary superpositions within this subspace. In this example, an arbitrary complex two-electron state is specified by $2\binom{4}{2} - 2 = 10$ real parameters. Since each A gate introduces two parameters, one might expect that only five A gates would be needed to generate an arbitrary state. However, we find that at least six A gates are needed to do this, although two of the parameters can be fixed to reduce the total parameter count back down to 10. We can again impose time-reversal symmetry by setting all the ϕ_i parameters to zero in the A gates, which reduces the parameter count down to 6—one more than the minimal number of 5. To remove this extra parameter, we find numerically that it works to set either θ_4 or θ_5 to zero. Note that setting both θ_i and ϕ_i to zero does not remove the i th A gate completely but instead reduces it to a CZ gate, as is evident from Eq. (2).

To confirm that this and all other circuits presented in this section indeed span the target subspace, we compute the fidelity $F = \frac{1}{N} \sum_{i=1}^N |\langle \Psi_i | \psi_i \rangle|^2$, where $|\Psi_i\rangle$ is a random state within the chosen subspace, and $|\psi_i\rangle$ is the state output by our circuit after we maximize $|\langle \Psi_i | \psi_i \rangle|^2$ with respect to the variational parameters. We check numerically that $F = 1$ can be achieved using the minimal number of parameters, $2\dim(H_{n,m}) - 2$. We choose enough random states, $N \gg \dim(H_{n,m})$, to ensure that the subspace is adequately represented. This is how we determine that, in the case of Fig. 3 with time-reversal symmetry, it works to set θ_4 or θ_5 to zero, but setting other θ parameters to zero instead does not achieve unit fidelity (although the fidelity still remains very high). Further details about the numerical verification of our circuits are given in the supplementary information.³¹

Although we have found a circuit that prepares all states in the two-excitation subspace using the minimal number of variational parameters (Fig. 3), this solution is neither unique nor optimal in terms of the number of CNOT gates. To illustrate these points, we present another circuit that accomplishes the same task in Fig. 4. In addition to the A gate, this circuit also makes use of single-qubit gates beyond just X gates. This example is quite different from the one shown in Fig. 3 in that it does not first apply X gates to two




Fig. 4 Alternate 4-qubit circuit. Another example circuit for the case of $n = 4$, $m = 2$, which also spans the desired subspace with the minimal 10 parameters, but with fewer two-qubit gates than the circuit shown in Fig. 3. The single-qubit gates $R(\theta, \phi)$ are as defined in Fig. 2.




Fig. 5 General A circuit. General construction of an efficient circuit which also enforces number symmetry for any number of qubits (orbitals) n and excitations (electrons) m and is constructed using the logic discussed in the text. Each A gate contributes two variational parameters θ_i, ϕ_i to the ansatz, except the last two A gates, which each contribute one (see text). This general structure only requires single-qubit X gates and a cascade of two-qubit A gates and always generates circuits with the minimal number of required parameters $2\binom{n}{m} - 2$.

qubits in order to bring the quantum processor into the appropriate particle-number subspace. Instead, the appropriate subspace is approached gradually as the circuit is performed, making it more challenging to understand and generalize the circuit. Since each A requires three CNOT gates to implement, we see that the circuit in Fig. 4 requires only 9 CNOT gates, while the one in Fig. 3 requires twice as many. However, it should be noted that the circuit of Fig. 3 requires only nearest-neighbor qubit coupling, which is not true of the one in Fig. 4. This example highlights the fact that further reductions in circuit depth are possible even if the circuit contains the minimal number of parameters, although this may require an increase in the qubit connectivity.

An efficient circuit for the general case of n orbitals and m fermions is shown in Fig. 5. This circuit builds on the approach of Fig. 3, where X gates are applied to m qubits to bring the system into the correct particle-number subspace, and then a sequence of A gates is performed to create different superpositions. Through trial and error, we find that the total number of A gates needed to produce all possible superpositions is $\binom{n}{m}$. Since each A gate contributes two parameters, we have a total of $2\binom{n}{m}$ parameters; the ϕ parameters in the last two A gates can be fixed to reduce the number of parameters down to the minimal number, $2\binom{n}{m} - 2$, while still fully spanning the subspace. Our general recipe for constructing circuits for any n, m like that shown in Fig. 5 can be summarized into the following steps:

1. Apply X gates to m qubits. For an efficient circuit, avoid placing X gates on neighboring qubits.
2. Apply a “first layer” of A gates on all adjacent pairs of qubits on which either $X \otimes 1$ or $1 \otimes X$ has been applied.
3. Apply a “second layer” of A gates on adjacent pairs of qubits, where each pair includes one qubit acted on by an A gate from the previous step and a qubit free of A gates. Continue

to place A gates on adjacent qubits as necessary until all neighboring qubits are connected with A gates. The first and second layers define a primitive pattern.

- Repeat the primitive pattern until $\binom{n}{m}$ A gates are placed. Any two A gates have ϕ as a free parameter and therefore the full circuit contains exactly $2\binom{n}{m} - 2$ parameters, the minimum required to span the subspace $H_{n,m}$.

In Step 1 above, we can see that this step simply places the full system into the proper m excitation subspace. The following two steps, which only involve applications of A gates, do not change the excitation number of the system but instead mix the $|01\rangle$ and $|10\rangle$ basis states of the two qubits on which they act, producing the desired subspace spanned by m , dependent on the parameters $(\vec{\theta}, \vec{\phi})$. Although we do not have an analytical argument for why this particular arrangement of A gates works, we have confirmed this through extensive numerical testing. We conjecture that this circuit pattern will continue to work for arbitrarily many fermions and orbitals. Another key point is contained in Step 4, specifically that this general construction always uses the minimal number of required parameters to span the desired space; increasing the gate depth any further is unnecessary, as unit fidelity is achievable at this gate depth. Time-reversal symmetry can be imposed by setting all the ϕ parameters and one θ parameter in the A gates to zero as before. This yields a purely real ansatz with the minimal number of parameters needed. Numerical calculations of the fidelity as a function of the number of variational parameters for several different values of n and m are given in Fig. S4 of the supplement.³¹ We have numerically verified all cases for which (n) and all permutations of $n = \{2, 3, 4, 5, 6\}$, $m = \{1, 2, 3, 4, 5\}$. In all cases, the fidelity increases monotonically with the number of parameters and saturates at unity once this number equals the dimension of the symmetry subspace, confirming that the circuits fully span the subspace.

While this construction is not necessarily the most resource efficient in terms of the required number of gates, it is straightforward to extend to any desired subspace defined by n and m and only requires nearest-neighbor qubit coupling, which is typically more straightforward to engineer. This should be contrasted with the circuit of Fig. 4, where two-qubit gates between non-neighboring qubits are required. We also note that our general construction naturally exhibits the symmetries of binomial functions. For example, the number of gates used for the case of $\binom{n}{p}$ and $\binom{n}{n-p}$ for $0 < p < n$ are identical, which is a reflection of particle-hole symmetry. This in turn means that we can focus on the case where $m \leq n/2$ without loss of generality, so that in Step 1 above, we can always avoid applying X gates on two adjacent qubits.

It is worth comparing our general construction with existing state preparation algorithms in terms of gate counts. If we wished to span the full Hilbert space, then this would require $\mathcal{O}(2^n)$ CNOT gates.^{32,33} Some previous state preparation algorithms involve transforming one arbitrary n -qubit state into another arbitrary state, which requires $2^{n+1} - 2n - 2$ CNOT gates.³⁴ However, since we are interested in only spanning a subspace of the full Hilbert space, namely $H_{n,m}$, we can span this subspace with significantly fewer CNOT gates. Wang et al.²⁴ and Ortiz et al.³⁵ also considered this restricted subspace, requiring no more than $2^{m+1}n^m/m!$ and $\binom{n}{m}^2 n^2$ CNOT gates, respectively. In the case of our general construction for n qubits and m excitations, we find that our algorithm requires at most $N_{\text{CNOT}}(n, m) = 3\binom{n}{m}$ CNOT gates. Since we always consider a fixed input state (the state with all qubits in $|0\rangle$), simplifications of our required gates are always possible, which reduces the number of required CNOT gates. Specifically, if we eliminate the unnecessary CNOT gates, then the actual




Fig. 6 CNOT scaling. Number of required CNOT gates as a function of particle number m for a fixed number of qubits $n = 40$. Our general state preparation circuits that respect particle-number symmetry are shown in blue, while the results of Wang et al.²⁴ are shown in yellow, Bergholm et al.³⁴ in green and Ortiz et al.³⁵ in red.

number of CNOT gates in the general circuit is

$$N_{\text{CNOT}}(n, m) = \begin{cases} 3\binom{n}{m} - 3m + 1 & 0/2 \\ 3\binom{n}{m} - 2m - 2 & m = n/2 \\ 3\binom{n}{m} - 3n + 3m + 1 & n/2. \end{cases} \quad (3)$$

Figure 6 shows how our approach to constructing state preparation circuits compares to existing works. We see that our scheme significantly decreases the required number of CNOT gates.

We noted earlier that our example of Fig. 4 contains only 9 CNOT gates, while our general construction uses $N_{\text{CNOT}}(4, 2) = 18$ CNOT gates in this case. However, our general construction in terms of A gates is relatively straightforward for arbitrary $H_{n,m}$, while it is not clear how to generalize the construction of Fig. 4. We stress that on a case by case basis, it may be possible to find circuits that more efficiently span a reduced Hilbert space, but finding a general procedure for constructing such circuits for arbitrary numbers of orbitals and electrons can be challenging, and it may require a more complicated qubit connectivity beyond just nearest-neighbor coupling. Further investigation of these differences and tradeoffs is an interesting topic for future work but beyond the scope of this paper.

Spin symmetries

Many systems that we are interested in simulating possess not only particle-number and time-reversal symmetries, but also spin symmetries. This can include both the net spin magnetization s_z and also total spin s . To our knowledge, it remains an open problem to find state preparation circuits that respect all these symmetries at the same time. Here, we introduce a protocol for constructing circuits that achieve this.

We begin by first showing how our general scheme for conserving particle number described in the previous section can be extended to conserve s_z as well with only minor modifications. The first step is to choose our fermion-qubit mapping such that the first $n/2$ qubits represent spin-up orbitals while the remaining $n/2$ represent spin-down orbitals. We follow the same general steps of forming a cascade of A gates as before, but now with the added constraint that the parameters of any A gate that bridges the two spin subspaces (those that entangle the $n/2$ qubit with the $n/2 + 1$ qubit) are set to zero ($\theta_i = \phi_i = 0$). This prevents any mixing of the two spin subspaces. Therefore, if we start with the proper number of spin-up and spin-down orbitals occupied, then this sequence of gates guarantees that the final state also has the correct spin occupation numbers. We also require that these parameter-free A gates do not appear in the

first layer. This can effectively swap the first and second layers as outlined in the previous general protocol. For example, in the case of Fig. 3 and with our spin assignment, we only need to set the four parameters contained within the A gates that act between qubits 2, 3 to zero along with $\phi_5 = \phi_4$. This assignment then generates all states with $n = 4, m = 2, s_z = 0$ and with the minimal number of parameters (6) for this space.

We now move on to the case where total spin s is also a good quantum number. If we impose conservation of s, s_z and particle number, then the size of the relevant symmetry subspace for a chosen n, m, s, s_z is given by

$$\dim(H_{n,m,s,s_z}) = \sum_{k=0}^{m/2-s} \binom{n/2}{k} \binom{n/2-k}{m-2k} \times \frac{(2s+1)(m-2k)!}{(m/2-k-s)!(m/2-k+s+1)!}. \quad (4)$$

Here, we maintain the notation from the previous section where n is the number of qubits and m is the number of particles, but now we focus on the case in which each qubit encodes the occupancy of a particular spin-orbital. (In the previous section each qubit could correspond to either a spin-orbital or a spatial orbital in the case of spinless fermions.) Therefore, the n qubits encode $n/2$ spatial orbitals, and we omit the case of odd values for n . The first binomial coefficient in Eq. (4) counts the number of ways to assign doubly occupied spatial orbitals in the system, the second counts the number of ways to assign singly occupied orbitals, and the last factor is the number of s irreducible representations in a tensor product of m spin- $\frac{1}{2}$ particles. We show in Fig. 7 that exploiting all the symmetries may significantly reduce the dimension of the Hilbert space. For example, already for $n = 28$ spin-orbitals, the relevant subspace when all symmetries are imposed is at least two orders of magnitude smaller than the full Hilbert space of 28 qubits. Taking advantage of this reduction can significantly reduce the demands on the quantum processor and improve the speed and accuracy of the classical optimization step of the VQE algorithm. We note that, while this reduction is significant, this scaling always remains exponential even when symmetries are imposed. Nevertheless, use of symmetries can simplify ansätze and thus reduce the required CNOT and parameter counts.

The problem of constructing circuits that also conserve total spin s is much more challenging compared to the symmetries we have discussed thus far. In the remainder of this section, we address this problem by developing a completely different approach to obtaining symmetry-preserving circuits. In addition




Fig. 7 Hilbert space scaling. Hilbert space dimension as a function of number of qubits, n , when relevant symmetries are enforced. We show the dimension of the full Hilbert space (blue), the largest particle-number subspace with $m = n/2$ (yellow), the subspace with $m = n/2$ and $s_z = 0$ (red), and the subspace with $m = n/2$ and $s = 0$ (green). Lines are included only as a guide. All cases remain an exponential scaling Hilbert space, where use of symmetries reduces the exponential factor.

to allowing for conservation of s , this approach also provides an alternative method for building circuits that respect particle-number symmetry. In this new approach, we adopt the convention that the qubits alternate between spin-up and spin-down orbitals such that the first, third, fifth, etc. qubits encode spin up, while the second, fourth, sixth, etc. encode spin down. This choice of spin encoding leads to an efficient scaling in terms of number of gates, as we will see.

The starting point for this alternative circuit-building scheme is to find an n -qubit unitary that transforms the state $\otimes_{i=1}^n |0\rangle$ into an arbitrary superposition of states that share the same value of m . This is of course exactly the problem of finding a circuit that respects particle-number symmetry that we discussed extensively in the previous section. The key idea here is that by starting with the explicit n -qubit unitary upfront instead of a decomposition of it in terms of A gates, we can control how circuit parameters appear in the coefficients of the final superposition states. If the circuit parameters appear in a sufficiently simple way, then it would be straightforward to impose constraints on them so that spin symmetries are also respected. We could attempt to use the particle-number circuits from the previous section and impose constraints on the circuit parameters to enforce total spin symmetry, but this would lead to complicated, highly non-linear conditions on the parameters that would be difficult to solve.

It is straightforward to construct an n -qubit unitary that transforms $\otimes_{i=1}^n |0\rangle$ into an arbitrary state in a particular particle-number subspace. In the simplest case of only two spin-orbitals and a single fermion, this unitary is just the A gate multiplied by a single X gate (see Fig. 1). In the case of four spin-orbitals and two fermions, the desired unitary should generate arbitrary superposition states formed from a total of $\binom{4}{2} = 6$ basis states. While many different unitaries transform $\otimes_{i=1}^n |0\rangle$ into such superposition states, here we desire a unitary which generates these states with *specific* coefficients that are easily adjusted to respect the appropriate spin symmetries. The unitary we choose to use, referred to as the E_4 gate, is based on hyperspherical coordinates:

$$\begin{aligned} E_4|0000\rangle &= \sin u_1 \sin u_2 \sin u_3 \sin u_4 \sin u_5 |0101\rangle \\ &+ \sin u_1 \sin u_2 \sin u_3 \sin u_4 \cos u_5 |1001\rangle \\ &+ \sin u_1 \sin u_2 \sin u_3 \cos u_4 |0011\rangle \\ &+ \sin u_1 \sin u_2 \cos u_3 |0110\rangle \\ &+ \sin u_1 \cos u_2 |1010\rangle \\ &+ \cos u_1 |1100\rangle. \end{aligned} \quad (5)$$

The E_4 gate clearly generates any two-particle state in a system with four spin-orbitals. Notice that we have also imposed time-reversal symmetry by purposely choosing the coefficients to be real. This can be easily generalized to problems without time-reversal symmetry by inserting additional arbitrary phase factors on any five of the six terms in Eq. (5). The use of hyperspherical coordinates provides a simple parameterization that automatically ensures normalization and facilitates the incorporation of spin symmetries. For example, if we want to restrict to the subspace with $s = 1, s_z = 0$, then there is only one state:

$$|s = 1, s_z = 0\rangle = \frac{1}{\sqrt{2}} (|1001\rangle + |0110\rangle).$$

This state can be created from the E_4 gate by setting $u_1 = u_2 = u_4 = \pi/2, u_3 = \pi/4, u_5 = 0$. We summarize this case and the other spin subspace of interest $s = 0, s_z = 0$, which is given by the general superposition

$$|s = 0, s_z = 0\rangle = \frac{\gamma}{\sqrt{2}} (|1001\rangle - |0110\rangle) + \alpha |1100\rangle + \beta |0011\rangle, \quad (6)$$

in Table 1.

The advantage of this construction is that we can easily fix parameters in the E_4 gate to generate any desired spin subspace in this $n = 4, m = 2$ space. There are of course many unitaries that satisfy Eq. (5). Once we have settled on a particular choice for E_4 , we can then decompose it into more elementary gates. Note that in Eq. (5) we have chosen a basis ordering which leads to an efficient Gray code decomposition³⁶ (there are only two bit changes from term to term). A specific choice for E_4 and its corresponding decomposition are shown in the appendix.

The case of $n > 4$ spin-orbitals can be treated in a similar fashion. We begin by constructing a unitary E_n that transforms $\otimes_{i=1}^n |0\rangle$ into an arbitrary superposition of $\binom{n}{m}$ fixed-particle-number basis states. For example, in the case $n = 6, m = 3$, we have

$$E_6|000000\rangle = \sum_{j=20}^1 \cos u_j \prod_{i=1}^{j-1} \sin u_i |p_{j1}, p_{j2}, \dots, p_{j6}\rangle, \quad (7)$$

where $p_{jk} \in \{0, 1\}$, $\sum_{k=1}^6 p_{jk} = 3$, $u_{20} = 0$, and the order of the basis states is given by

$$\begin{aligned} &|000111\rangle, |001011\rangle, |001110\rangle, |001101\rangle, |011001\rangle, |011010\rangle \\ &|011100\rangle, |010110\rangle, |010101\rangle, |010011\rangle, |100011\rangle, |100101\rangle \\ &|100110\rangle, |101010\rangle, |101100\rangle, |101001\rangle, |110001\rangle, |110010\rangle \\ &|110100\rangle, |111000\rangle. \end{aligned} \quad (8)$$

This ordering results in minimal (2) bit changes per step of the Gray code, which facilitates the decomposition of E_6 . We list the specific spin subspaces and summarize how to generate them by fixing the parameters of E_6 in Table 2.

An example of a gate decomposition for a specific choice of E_6 is given in the appendix along with explicit gate counts for E_4 , E_6 , and E_8 .

Extrapolating from these examples, we can then form a general procedure for constructing symmetry-preserving state preparation circuits for any valid choice of the quantum numbers n, m, s, s_z for a time-reversal-symmetric system as follows:

1. For a given choice of n, m , there are $\binom{n}{m}$ basis states $(|p_{j1}, p_{j2}, \dots, p_{jn}\rangle)$ that span the corresponding particle-

number subspace, such that $\sum_{k=1}^n p_{jk} = m$. Assign hyperspherical coefficients to these basis states according to $\sum_{j=\binom{n}{m}}^1 \cos u_j \prod_{i=1}^{j-1} \sin u_i |p_{j1}, p_{j2}, \dots, p_{jn}\rangle$, where $u_{\binom{n}{m}} = 0$, and the basis states are ordered such that there are only two bit changes per step in the Gray code.

2. Determine the constraints that must be imposed on the u_i such that the appropriate spin subspace labeled by s and s_z is obtained.
3. Construct a unitary E_n such that the first column contains the coefficients in Step 1. The remaining columns can be chosen as desired so long as they respect unitarity.
4. Decompose E_n using a Gray code scheme or alternative gate decomposition technique.

Following this procedure, we can enforce arbitrary spin and particle-number symmetries for any choice of n, m, s, s_z . Step 2 is facilitated by the use of hyperspherical coordinates because various basis states are easily eliminated from the final superposition by setting the corresponding u_j to $\pi/2$, yielding another hypersphere parameterization of lower dimension. This recursive structure in the coefficients maintains regularity and simplicity as the effective Hilbert space dimension is reduced. Note that in cases where it is necessary to fix some of the u_j in terms of the others (see e.g., the case of $s = s_z = 0$ in Table 1 or $s = 1/2 = \pm s_z$ in Table 2), use of inverse functions without a restricted real-domain (e.g. \tan^{-1}, \cot^{-1}) are desirable for optimization stability and enforcement of time-reversal symmetry, while the use of other inverse trigonometric functions could yield complex-valued coefficients in the trial states. Finally, we emphasize that here it is not necessary to numerically compute fidelities to confirm that the resulting circuits indeed span the appropriate symmetry subspace as in the previous section because this is guaranteed by construction in the present approach.

To demonstrate the efficiency of our state preparation circuits, we use them in a VQE simulation to compute the ground state energy of the H_2 molecule. We work in the STO-3G basis and map to qubits using the Jordan-Wigner transform.

We note that there are many choices of qubit mappings and unitary transformations that can be used to reduce the complexity of molecular simulations.³⁷ However, we have formulated our ansätze to preserve symmetries in the natural basis where each qubit represents a physical spin orbital. There are also several techniques to tailor an ansatz to efficiently describe the allowed transitions in H_2 , but here we are concerned with constructing general ansätze, applicable to any molecule, which only utilize particle number and spin symmetries, to reduce the complexity of the resulting circuits. Incorporation of particle number, spin projection and total spin symmetries in conjunction with other qubit mappings and reduction methods is an interesting topic, but beyond the scope of the present work. The results are given in Fig. 8, which shows the difference between the computed ground state energy and the exact ground state energy. For comparison, we also show the results obtained from standard ansätze (SWAPRZ, RY, RYRZ, and UCCSD) included in IBM's Qiskit software package.³⁸

Table 1. Two particle spin symmetry.

s, s_z	u_1	u_2	u_3	u_4	u_5
1, 0	$\pi/2$	$\pi/2$	$\pi/4$	$\pi/2$	0
0, 0	u_1	$\pi/2$	$-\tan^{-1}(\csc u_4)$	u_4	0

By fixing the coefficients of Eq. 5, we can generate the two spin subspaces defined by their quantum numbers for the $n = 4, m = 2$ space. In the case of $s = 1, s_z = 0$, the subspace is spanned by a single state, so all the u_i are fixed. The subspace with $s = s_z = 0$ is three-dimensional, and so two of the u_i are left unspecified

Table 2. Three particle Spin Symmetry.

s, s_z	u_1	u_2	u_3	u_4	u_5	u_6	u_7	u_8	u_9	u_{10}	u_{11}	u_{12}	u_{13}	u_{14}	u_{15}	u_{16}	u_{17}	u_{18}	u_{19}
$\frac{3}{2}, \frac{1}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\cot^{-1}(\frac{1}{\sqrt{2}})$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{4}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	0	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$
$\frac{3}{2}, -\frac{1}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\cot^{-1}(\frac{1}{\sqrt{2}})$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{4}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	0	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$
$\frac{1}{2}, \frac{1}{2}$	u_1	$\frac{\pi}{2}$	u_3	$\frac{\pi}{2}$	$-\cot^{-1}(\delta)$	u_6	$\frac{\pi}{2}$	u_8	$\frac{\pi}{2}$	u_{10}	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	u_{15}	$\frac{\pi}{2}$	$\frac{\pi}{2}$	u_{18}	0
$\frac{1}{2}, -\frac{1}{2}$	$\frac{\pi}{2}$	u_2	$\frac{\pi}{2}$	u_4	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$-\cot^{-1}(\kappa)$	$\frac{\pi}{2}$	u_{11}	$\frac{\pi}{2}$	u_{13}	u_{14}	$\frac{\pi}{2}$	u_{16}	u_{17}	$\frac{\pi}{2}$	$\frac{\pi}{2}$

By fixing the coefficients of Eq. 7, we can generate the four spin subspaces defined by their spin quantum numbers for the $n = 6, m = 3$ space. Here $\delta = \{\sin(u_{10}) \cos(u_{15}) \sin(u_6) \sin(u_8) + \sin(u_6) \cos(u_8)\}$, $\kappa = \{\sin(u_{11}) \sin(u_{13}) \sin(u_{14}) \cos(u_{16}) + \sin(u_{11}) \cos(u_{13})\}$. The order of basis states for this space is defined in the text

Furthermore, we show the variational parameter and CNOT gate counts in Table 3, accompanying Fig. 8 for all ansätze considered. It is clear that our proposed circuits obtain the ground state energy with high accuracy using the minimal number of variational parameters and a low number of CNOT gates. For the RY, RYRZ, and SWAPRZ ansätze, we choose a circuit depth of three, while for UCCSD, we use a circuit depth of one. In all cases, we use the BFGS optimizer (run several times with random initial conditions) included in Qiskit and consider an ideal, noiseless simulation.

Figure 9 shows the results of a similar simulation, but this time with noise included. This noise is characterized by the physical parameters of IBM's Poughkeepsie device. Here we only take the standard ansatz that has reasonable parameter and CNOT scaling, RY, and compare it to our ansätze. Contrary to the noiseless case, the RY ansatz here is chosen to have linear connectivity (nearest-neighbor CNOT connectivity) and a circuit depth of one as this configuration performs best under these noisy conditions. With these settings, the RY ansatz has 8 parameters and 3 CNOT gates. The CNOT and parameter counts for our ansätze are the same as in Fig. 8 since these are determined by the ground state symmetries of the molecule. Even though the RY ansatz has fewer costly CNOT gates than our A gate ansatz and is overparameterized for the target ground state, it is outperformed by the A gate ansatz. This is due to the fact that at this low depth, the RY ansatz does not have sufficient coverage of the target Hilbert space. Increasing the depth of the RY ansatz will increase its




Fig. 8 H_2 Noiseless VQE. Energy difference from the exact ground state of H_2 for various ansätze. All ansätze perform well below chemical accuracy, but vary significantly in their number of required parameters and CNOT gates (shown in the accompanying table). Our proposed ansätze ($A_{4,2}$ and $E_{4,2}$) achieve very small energy differences and have low to modest CNOT counts. The UCCSD ansatz also performs very well with low parameters (3), but has a fairly large CNOT count in this example.

Table 3. 4-qubit ansätze comparison.

Ansätze	# CNOT	# Parameter
$A_{4,2,s_z=0}$	6	3
$E_{4,2}$	14	5
$E_{4,2,s=0,s_z=0}$	20	2
UCCSD	56	3
SWAPRZ	34	72
RYRZ	18	32
RY	18	16

In the ideal, noiseless case, we compare many standard ansätze to our new proposed ansätze and find that our proposed ansätze are more compact in terms of CNOT count and parameter count

coverage of the target space, but also increases the number of CNOT gates. This is not an equivalent trade-off and actually decreases the ability of the ansatz to approximate the ground state. In the case of both E gate ansätze, their relatively larger number of CNOT gates can be seen to hinder their performance. However, as these gates are constructed in terms of arbitrary state generation, these gate counts could be improved with more efficient state generation algorithms. In addition, mitigation of CNOT errors may be possible with the use of Richardson extrapolation,^{39,40} a technique we will pursue in future work.

An interesting feature in Fig. 9a) can be seen where the $E_{4,2}$ ansatz crosses the $E_{4,2,s=0,s_z=0}$ ansatz. In principle, the spin symmetry ansatz should always have a better ability to target the ground state in a VQE algorithm. This is explained in Fig. 9b), where as we see that the $E_{4,2}$ ansatz abruptly changes the total spin value of its lowest energy state. Indeed, the triplet states ($s = 1$) are nearly degenerate with the singlet state ($s = 0$, Eq. 6) at large bond distances. At an interatomic distance near two Angstroms, the energy difference between the lowest energy singlet and triplet states is $\mathcal{O}(10^{-2})$ Hartree. The optimization landscape becomes much more difficult to navigate when there are many nearly equal local minima, as is the case here at large bond distances. Therefore since the algorithm attempts to find the lowest energy eigenvalue and the singlet and triplet states are nearly degenerate at longer bond distances, the optimizer will frequently return a minimum (triplet state) rather than the global minimum (singlet state). In the presence of noise, this issue is blurred further as both energies are not exactly attainable, but a triplet state is trivial to produce (e.g. $s = 1, s_z = 1, |1010\rangle$), so finding the triplet state as the lowest energy state is frequently encountered. Additionally, this issue grows more difficult to navigate with larger molecules and strongly correlated systems, which have many more nearly degenerate ground states. Naturally one could also enforce a spin eigenvalue constraint on the optimizer by modifying the objective function, but this places stricter requirements on the optimizer and requires a physical measurement of spin eigenvalues, through tomography or additional circuit elements. However, enforcing symmetry in the ansätze lessens the burden on the optimizer and does not directly require spin measurements.

We also perform similar noisy simulations on LiH . LiH in the STO-3G basis maps to 12 qubits but we follow the same methods as ref.,⁸ removing two non-interacting orbitals and freezing the core orbital, reducing LiH to 6 qubits. We show in Fig. 10 and Table 4 that our ansätze perform comparably for LiH as they do for H_2 with notably more error for the E gate ansatz as it requires significantly more CNOTs. As in the noisy H_2 case, the RY ansatz is chosen to have a fixed depth (one) and connectivity (nearest-neighbor). These settings give the best performance under these noisy conditions but the same coverage-CNOT trade off issue arises as before. From both the H_2 and LiH results we can see that the A -gate circuit finds a middle-ground in terms of CNOT and parameter count and has good performance, while the E -gate circuit exactly minimizes parameter count based on particle number, spin projection, and total spin, but its performance suffers due to a large CNOT count. Thus, in the presence of noise, we find that it can be favorable to relax some of the symmetry constraints to reduce the circuit depth at the expense of more variational parameters.

Since our constructions maintain number and spin symmetries, potentially any noise sources which violate these symmetries can be mitigated through post-selection or symmetry verification.^{13,14,27} This will be a topic of future work. Another interesting question for future work is whether it is possible to find circuits that minimize the number of CNOT gates while respecting spin symmetries and while using the minimal number of parameters necessary to span the symmetry subspace.

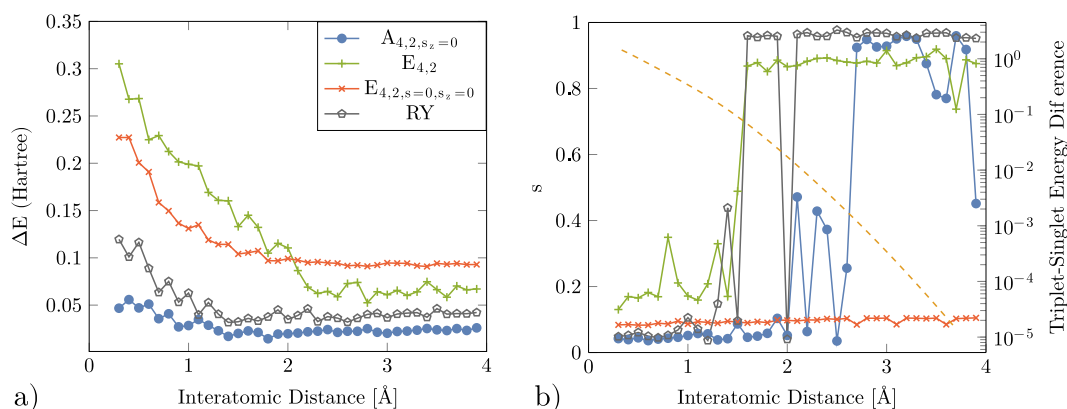


Fig. 9 H_2 Noisy VQE. **a)** Energy difference from the exact ground state of H_2 in the presence of noise characterized by IBM's Poughkeepsie device. Included is Qiskit's state preparation and measurement error mitigation for all ansätze. **b)** Total spin eigenvalues for the same dissociation curve (left axis). The total spin values indicate that some ansätze drift outside the singlet subspace ($s = 0$) and begin to instead find a (nearly degenerate) triplet state ($s = 1$) at large bond distances. Since the $E_{4,2,s=0,s_z=0}$ ansatz is restricted to a fixed total spin, it is the only displayed ansätze which always attempts to find the true ground state across the full dissociation curve. The energy difference between this triplet state and the true singlet ground state as a function of interatomic distance is also shown for comparison (orange dashed line, right axis).




Fig. 10 LiH Noisy VQE. Energy difference from the exact ground state of LiH in the presence of noise characterized by IBM's Poughkeepsie device. With appropriate orbital reductions, LiH is run on 6 qubits and has two excitations. The A-gate ansatz continues to perform well utilizing particle number and spin projection symmetries and relatively low CNOT count for this dimension.

Table 4. 6-qubit ansätze comparison.

Ansätze	# CNOT	# Parameter
$A_{6,2,s_z=0}$	25	8
$E_{6,2,s=0,s_z=0}$	48	5
RY	5	12

CNOT and parameter count for the case of estimating the ground state of LiH in the presence of typical device noise

DISCUSSION

In this work, we presented general schemes to construct state preparation circuits for quantum simulation that respect a number of symmetries that commonly arise in systems of interest, including particle number, time-reversal, total spin, and spin magnetization. In each case, we provide general construction procedures, explicit examples of circuits, and gate counts. In the case of particle-number symmetry, for which state preparation circuits have been

proposed previously by other authors, our circuits outperform existing methods in terms of the number of two-qubit entangling gates they contain. Enforcing spin symmetries in addition to particle number can significantly enhance the performance of the VQE by preventing the algorithm from wasting time exploring vast, irrelevant regions of Hilbert space. Also, our ansätze are *guaranteed* to contain the ground state of interest, with the minimal number of parameters defined by the symmetry subspace. It is important to emphasize, however, that the number of variational parameters still grow exponentially with system size even when all symmetries are imposed. Thus, spanning the entire symmetry subspace will become impractical as the size of the quantum processor is increased beyond a few tens of qubits. However, we expect that our symmetry preservation techniques will continue to play an important role for larger NISQ devices as systematic methods to further reduce the search to still smaller regions of Hilbert space are developed. This could be done, for example, by looking for ways to freeze large sets of the variational parameters appearing in our circuits without sacrificing the accuracy of the ansatz. In addition, as a general state creation circuit with a particular symmetry, one can imagine applications for these circuits in the Phase Estimation Algorithm (PEA).

METHODS

Confirmation of circuit compilation was performed in Mathematica software as well as numerical optimization to confirm that the proposed circuits span any state in the desired Hilbert space. In numerical optimization testing, we find that Mathematica's NMaximize method "Simulated Annealing" performs best for our purposes, when run over many samples of a randomly chosen state in the Hilbert space. Gate decomposition was performed using the method of Gray codes, was done "by hand" and confirmed in Mathematica.

Our VQE simulations were done using IBM's Qiskit software, coding in our own proposed ansätze, which enforce relevant symmetries. Noisy simulations were simulated using the noise parameters of IBM's Poughkeepsie device.

DATA AVAILABILITY

The data that support the findings of this study are available from the authors upon request.

CODE AVAILABILITY

A custom Mathematica code to reproduce our results is available on GitHub.⁴¹

Received: 25 April 2019; Accepted: 10 November 2019;
Published online: 28 January 2020

REFERENCES

- Feynman, R. P. Simulating physics with computers. *Int. J. Theoretical Phys.* **21**, 467–488 (1982).
- Abrams, D. S. & Lloyd, S. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.* **83**, 5162–5165 (1999).
- Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **309**, 1704–1707 (2005).
- Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
- Kühn, M., Zanker, S., Deglmann, P., Marthaler, M. & Weiß, H. Accuracy and resource estimations for quantum chemistry on a near-term quantum computer. *J. Chem. Theory Comput.* **15**, 4764 (2019).
- Peruzzo, A. et al. A variational eigenvalue solver on a quantum processor. *Nat. Commun.* **5**, 4213 (2014).
- O'Malley, P. J. J. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
- Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242 (2017).
- Colless, J. I. et al. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Phys. Rev. X* **8**, 011021 (2018).
- Shen, Y. et al. Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure. *Phys. Rev. A* **95**, 020501 (2017).
- Hempel, C. et al. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X* **8**, 031022 (2018).
- Nam, Y. et al. Ground-state energy estimation of the water molecule on a trapped ion quantum computer. Preprint at <https://arxiv.org/abs/1902.10171> (2019).
- McArdle, S., Yuan, X. & Benjamin, S. Error-mitigated digital quantum simulation. *Phys. Rev. Lett.* **122**, 180501 (2019).
- Bonet-Monroig, X., Sagastizabal, R., Singh, M. & O'Brien, T. E. Low-cost error mitigation by symmetry verification. *Phys. Rev. A* **98**, 062339 (2018).
- Barkoutsos, P. K. et al. Quantum algorithms for electronic structure calculations: particle-hole hamiltonian and optimized wave-function expansions. *Phys. Rev. A* **98**, 022322 (2018).
- Bartlett, R. J., Kucharski, S. A. & Noga, J. Alternative coupled-cluster ansätze ii. the unitary coupled-cluster method. *Chemical Physics Letters* **155**, 133–140 (1989).
- Kutzelnigg, W. Error analysis and improvements of coupled-cluster theory. *Theoretica Chimica Acta* **80**, 349–386 (1991).
- Ryabinkin, I. G., Yen, T.-C., Genin, S. N. & Izmaylov, A. F. Qubit coupled cluster method: a systematic approach to quantum chemistry on a quantum computer. *J. Chem. Theory Comput.* **14**, 6317–6326 (2018).
- Grimsley, H. R., Economou, S. E., Barnes, E. & Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* **10**, 3007 (2019).
- Bian, T., Murphy, D., Xia, R., Daskin, A. & Kais, S. Quantum computing methods for electronic states of the water molecule. *Mol. Phys.* **117**, 2069–2082 (2019).
- McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
- McClean, J. R., Romero, J., Babbush, R. & Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **18**, 023023 (2016).
- Ryabinkin, I. G., Genin, S. N. & Izmaylov, A. F. Constrained variational quantum eigensolver: quantum computer search engine in the fock space. *J. Chem. Theory Comput.* **15**, 249–255 (2019).
- Wang, H., Ashhab, S. & Nori, F. Efficient quantum algorithm for preparing molecular-system-like states on a quantum computer. *Phys. Rev. A* **79**, 042335 (2009).
- Roth, M. et al. Analysis of a parametrically driven exchange-type gate and a two-photon excitation gate between superconducting qubits. *Phys. Rev. A* **96**, 062323 (2017).
- Egger, D. et al. Entanglement generation in superconducting qubits using holonomic operations. *Phys. Rev. Applied* **11**, 014017 (2019).
- Sagastizabal, R. et al. Error mitigation by symmetry verification on a variational quantum eigensolver. *Phys. Rev. A* **100**, 010302 (2019).
- Ganzhorn, M. et al. Gate-efficient simulation of molecular eigenstates on a quantum computer. *Phys. Rev. Applied* **11**, 044092 (2019).
- Jordan, P. & Wigner, E. über das paulische äquivalenzverbot. *Z. Phys.* **47**, 631 (1928).
- Zhang, J., Vala, J., Sastry, S. & Whaley, K. B. Geometric theory of nonlocal two-qubit operations. *Phys. Rev. A* **67**, 042313 (2003).
- Gard, B., Zhu, L., Barron, G., Barnes, E. & Economou, S. Supplementary Information for: Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm. <https://doi.org/10.1038/s41534-019-0240-1>.
- Shende, V. V. & Markov, I. L. Quantum circuits for incompletely specified two-qubit operators. *Quantum Inf. Comput.* **5**, 49 (2005).
- Möttönen, M., Vartiainen, J. J., Bergholm, V. & Salomaa, M. M. Quantum circuits for general multiqubit gates. *Phys. Rev. Lett.* **93**, 130502 (2004).
- Bergholm, V., Vartiainen, J. J., Möttönen, M. & Salomaa, M. M. Quantum circuits with uniformly controlled one-qubit gates. *Phys. Rev. A* **71**, 052330 (2005).
- Ortiz, G., Gubernatis, J. E., Knill, E. & Laflamme, R. Quantum algorithms for fermionic simulations. *Phys. Rev. A* **64**, 022319 (2001).
- Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. (Cambridge University Press, New York, 2011).
- Bravyi, S., Gambetta, J.M., Mezzacapo, A. & Temme, K. Tapering off qubits to simulate fermionic hamiltonians. Preprint at <https://arxiv.org/abs/1905.01534> (2017).
- Abraham, H. et al. Qiskit: An open-source framework for quantum computing. <https://qiskit.org/>. (2019).
- Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017).
- McCaskey, A.J. et al. Quantum chemistry as a benchmark for near-term quantum computers. Preprint at <https://arxiv.org/abs/1905.01534> (2019).
- Gard, B., Zhu, L., Barron, G., Barnes, E. & Economou, S. <https://github.com/bgard/1/Efficient-Symmetry-Preserving-State-Preparation-Circuits-for-the-VQE-algorithm>. (2019).

ACKNOWLEDGEMENTS

This research was supported by the US Department of Energy (Award No. de-sc 0019199) and the National Science Foundation (Award No. 1839136). S.E.E. also acknowledges support from Award No. de-sc 0019318 from the Department of Energy. This research used quantum computing system resources supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research program office. Oak Ridge National Laboratory manages access to the IBM Q System as part of the IBM Q Network.

AUTHOR CONTRIBUTIONS

S.E.E., E.B., and N.J.M. structured and supervised the project. E.B. proposed methods for symmetry encoding and minimal parameterization. B.T.G. performed simulations, formulated general circuits and drafted the manuscript. L.Z. performed circuit decomposition, CNOT gate counts and suggested example circuits. G.S.B. drafted figures and confirmed circuit compilations. All authors contributed to the revisions of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information is available for this paper at <https://doi.org/10.1038/s41534-019-0240-1>.

Correspondence and requests for materials should be addressed to B.T.G.

Reprints and permission information is available at <http://www.nature.com/reprints>

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



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

© The Author(s) 2020