



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Progress towards practical quantum variational algorithms

Dave Wecker, Matthew B. Hastings, and Matthias Troyer

Phys. Rev. A **92**, 042303 — Published 2 October 2015

DOI: [10.1103/PhysRevA.92.042303](https://doi.org/10.1103/PhysRevA.92.042303)

Towards Practical Quantum Variational Algorithms

Dave Wecker,¹ Matthew B. Hastings,^{2,1} and Matthias Troyer³

¹*Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA*

²*Station Q, Microsoft Research, Santa Barbara, CA 93106-6105, USA*

³*Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland*

The preparation of quantum states using short quantum circuits is one of the most promising near-term applications of small quantum computers, especially if the circuit is short enough and the fidelity of gates high enough that it can be executed without quantum error correction. Such quantum state preparation can be used in variational approaches, optimizing parameters in the circuit to minimize the energy of the constructed quantum state for a given problem Hamiltonian. For this purpose we propose a simple-to-implement class of quantum states motivated by adiabatic state preparation. We test its accuracy and determine the required circuit depth for a Hubbard model on ladders with up to 12 sites (24 spin-orbitals), and for small molecules. We find that this ansatz converges faster than previously proposed schemes based on unitary coupled clusters. While the required number of measurements is astronomically large for quantum chemistry applications to molecules, applying the variational approach to the Hubbard model (and related models) is found to be far less demanding and potentially practical on small quantum computers. We also discuss another application of quantum state preparation using short quantum circuits, to prepare trial ground states of models faster than using adiabatic state preparation.

PACS numbers:

I. INTRODUCTION

Variational classes of states such as matrix product states (MPS) [1], multiscale entanglement renormalization (MERA) [2], and projected entangled pair states (PEPS) [3] play a key role in studying many-body quantum systems. An ideal class of states should be large enough to approximate the ground state and it must be possible to evaluate the energy and observables. Unfortunately for many potential classes of states, including PEPS, evaluation may be difficult on a classical computer, and for other classes it can be very computationally expensive. However, on a *quantum computer*, large classes of PEPS states can be prepared efficiently [4, 5], (although likely not all PEPS can be prepared this way [6]). Once the state is prepared, the energy and observables can be measured by sampling using short quantum circuits. Similar quantum circuits exist for MPS and MERA states [7]. Recently, Ref. 8 proposed variational methods for studying quantum chemistry problems on a quantum computer, and demonstrated one method on a model with a four-dimensional Hilbert space.

Preparation of variational states is an attractive application of small quantum computers, since many classically intractable states can be prepared with quite short quantum circuits and thus do not pose stringent requirements on coherence times and gate fidelities. While variation over all possible states that can be produced by quantum circuits of a given maximum depth (as done in the demonstration experiments of Ref. 8) might sound ideal, this approach does not scale efficiently. Ref. 8 also proposed using the unitary coupled cluster method (UCC) [9], where variational states are prepared using unitary evolution under a sum of fermionic terms including $c_p^\dagger c_q + h.c.$, $c_p^\dagger c_q^\dagger c_r c_s + h.c.$, and higher order terms,

typically restricted to the case that creation operators are on unoccupied orbitals in the Hartree-Fock state and annihilation operators are on occupied (or vice-versa as required by hermiticity). That approach involves many parameters: even if truncated at the level of four fermion operators, the number of variational parameters scales as the number of occupied orbitals squared times the number of unoccupied orbitals squared, which at constant filling fraction is the fourth power of the number of orbitals.

In this paper, we present an analysis of a different class of variational states, that we term “Hamiltonian variational”, using a modest gate depth and a very modest number of variational parameters compared to the system size. We present a detailed numerical analysis of this variational technique applied to a Hubbard model, ranging up to systems of 12 sites (equivalently, 24 spin-orbitals) at half-filling, to quantify its accuracy. The difficulty in general increases with system size, but also fluctuates in complicated ways, depending upon the spectrum of the non-interacting model. For certain cases, including our largest size, the non-interacting model is highly degenerate, leading to strong interaction effects and a very poor overlap of an initial Slater determinant state with the true ground state. We find that even then the variational approach works well. Finally, we also analyze the effects of sampling error in measuring energy and give concrete estimates for time scales to implement this procedure in practice.

The Hamiltonian variational method builds the variational state using rotations by terms in the Hamiltonian. This is then very well suited for models such as the Hubbard model where there are few interaction terms, all of which take a simple form; in contrast, for UCC applied to the Hubbard model, the large number of possible terms

$c_p^\dagger c_q^\dagger c_r c_s + h.c.$ will lead to a much larger circuit depth and further each term will take a more complicated form than the simple interaction in the Hubbard model. So, for comparison purposes, we instead apply both methods to several quantum chemistry systems also where the number of terms in the Hamiltonian is comparable to the number of possible four fermion terms. In this case, we consider additionally several variants of UCC that we term Rxx (different variants include different terms) described below and we show that it is possible to improve the UCC circuit depth by taking a very limited Trotter number. While UCC shows a high accuracy on small molecules, we find that this requires more evaluations. Further, on larger molecules with stronger interaction, we find that the accuracy of UCC and Rxx is worse than that obtained by the Hamiltonian variational method.

The class of variational states that we consider is inspired by both adiabatic state preparation and the quantum optimization algorithm of Refs. 10, 11. Consider a family of Hamiltonians $H = \sum_a J_a h_a$ with the J_a being scalars and the h_a being operators. Let J_a^0 and J_a^1 be two choices of these J_a , with corresponding Hamiltonians H_0, H_1 . Suppose it is easy to prepare a ground state Ψ_I of H_0 (for example, in our study of the Hubbard model, H_0 is non-interacting). Then, assuming that no gap closes, it is possible to adiabatically evolve from Ψ_I to the ground state of H_1 . If we break the annealing into short time steps dt and evolve for a total time T , this annealing is a sequence of (T/dt) different unitary rotations by Hamiltonians interpolating between H_0 and H_1 . This could be implemented on a quantum computer using a Trotter-Suzuki method which further decomposes this sequence into a sequence of unitary rotations by individual terms in the Hamiltonian.

The idea of the variational method that we consider is still restricted to a sequence of unitary rotations by terms in the Hamiltonian, but considers arbitrary angles for the rotations in the sequence, rather than choosing them from a Trotterization of an annealing process. By choosing these angles arbitrarily, this allows us to take a much shorter sequence. Thus, we consider a trial state, which we call the ‘‘Hamiltonian variational’’ state

$$\Psi_T = \exp(i\theta_n h_{a_n}) \dots \exp(i\theta_2 h_{a_2}) \exp(i\theta_1 h_{a_1}) \Psi_I, \quad (1)$$

where larger n increases the accuracy. The angles θ_k are variational parameters. The optimization algorithm of Refs. 10, 11 considered only rotation by two different types of operators, while we consider rotation by a larger number of terms.

If all the terms in the Hamiltonian have some symmetry (such as spin-rotation symmetry), then Ψ_I and Ψ_T transform in the same way under this symmetry. This allows us to find ground states with different quantum numbers by picking initial states Ψ_I with the desired quantum number. In cases of symmetries such as translation invariance, it is helpful to try to construct the sequence of terms a_k to approximately preserve this symmetry (see a more detailed discussion below).

To be useful for variational state preparation, it must be possible to optimize over the given parameters, without getting stuck in false minima. In Sec. IID we discuss techniques which succeed in finding good optima for the Hubbard model using access to numerically exact values of the energy obtained from a classical simulation. Such a classical approach may be useful on a quantum computer as it finds short circuits that prepare specific highly entangled states, albeit limited to system sizes where classical simulation is possible.

To go beyond the limits of classical simulation on a quantum computer one must measure the energy using a quantum circuit. One way is to write the Hamiltonian as a sum of sets of terms, so that all terms in a given set commute with each other and then measure each set in one run, doing many runs for each set of terms, with the error in the energy going as the inverse square-root of the number of samples. The required number of samples to achieve high accuracy may be large and limiting the number of samples significantly impacts the ability to find the optimum as discussed later in Sec. IIE.

Another way to measure energy is to use phase estimation to compute the energies, rather than sampling. In this case, the error in energy goes inversely with the phase estimation time. This procedure still is probabilistic, returning on each run an energy chosen randomly from a distribution. Thus many runs are still required to estimate the average energy. For many distributions, this procedure provides much more accurate access to the energy for a given run time than sampling would, at the cost of requiring longer coherence time. Additionally, it allows one to optimize on the probability of finding the ground state, rather than on the energy, which may improve the optimization.

One may wonder why this would be useful, since if phase estimation finds an energy equal to the ground state, one knows that one has successfully prepared the ground state and can now make a measurement, without any need to improve the variational state. However, once the optimum variational parameters are found, these parameters are classical information that can be used to quickly re-create the state. This is useful if one wishes to make many measurements on the state, if the measurements destroy the state. Also, techniques in Ref. 12 show how to quadratically speed up the sampling of properties of the state, assuming access by measuring a projector onto the state. However, if we have access to a projector onto Ψ_I , then conjugating this projector by the unitaries in Eq. (1) gives us a short-depth circuit which projects Ψ_T . Finally, in many applications such as chemical reactions where one studies the same Hamiltonian along a path of parameters, it may be possible to use the variational solution for a given Hamiltonian to find a solution of a nearby Hamiltonian rapidly.

II. THE HUBBARD MODEL

A. The Hubbard Hamiltonian

To study performance scaling across a range of sizes, we consider a sequence of Hubbard models on N -sites, with the sites arranged in a two-leg ladder (an $N/2$ -by-2 square lattice) for $N = 4, 6, 8, 10, 12$. The Hamiltonian of the Hubbard model is

$$H = -t \sum_{\langle i,j \rangle} \sum_{\sigma} t_{ij} c_{i,\sigma}^{\dagger} c_{j,\sigma} + U \sum_i c_{i,\uparrow}^{\dagger} c_{i,\uparrow} c_{i,\downarrow}^{\dagger} c_{i,\downarrow}, \quad (2)$$

with $t = 1, U = 2$. Here $c_{i,\sigma}^{\dagger}$ and $c_{i,\sigma}$ create and annihilate an electron at site i with spin σ respectively. We use periodic boundary conditions along the long “horizontal” direction of the lattice which is $N/2$ sites long and open boundary conditions in the short “vertical” direction so that there are a total of N bonds in the horizontal direction and $N/2$ in the vertical direction, all with the same strength. We write this as $H = h_h + h_v + h_U$, where h_h is the sum of hopping terms in the horizontal direction, h_v is the sum of hopping terms in the vertical direction, and h_U is the repulsion term.

B. Ground state degeneracies and initial states

We studied the system at half-filling with a N electrons on N sites. For an equal number of $N/2$ up and down electrons the single particle spectrum for $N = 4, 6, 10$ has $N/2 - 1$ states below the Fermi energy and is doubly degenerate at the Fermi energy, giving a $2^2 = 4$ -fold degenerate many-body ground state for the non-interacting ($U = 0$) Hamiltonian. For $N = 8$ the $U = 0$ Hamiltonian has a unique ground state, with $N/2$ states below the Fermi energy and an excitation gap. For $N = 12$, there are $N/2 - 2$ single particle states below the Fermi energy, and four states at the Fermi energy, so that for $N/2$ up electrons and $N/2$ down electrons, the $U = 0$ ground state is $\binom{4}{2}^2 = 36$ -fold degenerate. These different degeneracies of the non-interacting problem impact the difficulty of solving it. They also reduce the overlap of a Slater determinant with the true ground state. The degeneracies at $N = 4$ and 12 are due to the special choice of the vertical coupling; for generic vertical coupling, degeneracies are only seen for $N = 4n + 2 = 6, 10, 14, \dots$, where n is a positive integer.

The spin of the ground state at $U = 2$ also depends upon N . It is a singlet for $N = 4n = 4, 8, 12, \dots$ and is a triplet for $N = 4n + 2 = 6, 10, \dots$. All the results that we report below are for Ψ_I in the correct spin sector. In the case of $N = 6, 10$ we did this by choosing $N/2 + 1$ up electrons and $N/2 - 1$ down electrons. In that sector, the free fermion ground state becomes non-degenerate. Otherwise, we chose $N/2$ up electrons and $N/2$ down electrons. To choose a unique ground state for $N = 4, 12$ we prepared the ground state of the Hamiltonian $th_h +$

$(1 - \epsilon)th_v$ for small $\epsilon > 0$; this state is independent of ϵ for ϵ small and this simply picks out one of the ground states of H .

C. The variational ansatz

We choose the terms a_k in Eq. (1) in a repeating pattern and perform S repetitions, which we call “steps”. In each step, we have three variational parameters, $\theta_h^b, \theta_v^b, \theta_U^b$, where $b = 1, \dots, S$. We set

$$\Psi_T = \prod_{b=1}^S \left(U_U \left(\frac{\theta_U^b}{2} \right) U_h(\theta_h^b) U_v(\theta_v^b) U_U \left(\frac{\theta_U^b}{2} \right) \right) \Psi_I, \quad (3)$$

where $U_X(\theta)$ approximates $\exp(i\theta h_X)$ for $X \in \{U, h, v\}$ and the product is ordered by decreasing b ; we say “approximates” because for h_h is a sum of non-commuting terms and we therefore used a second-order Trotter-Suzuki method to implement U_h , applying the terms in sequence from left to right in each row and then reversing the order. This Trotterization helps to approximately preserve momentum, which is useful if Ψ_I has the right momentum quantum numbers. For U_U, U_v , we implement the exponential exactly. Each step of the ansatz then is a second-order Trotter approximation to evolution under $H(b) \equiv \theta_U^b h_U + \theta_v^b h_v + \theta_h^b h_h$, (note that $[U_h, U_v] = 0$; this property holds in general for free fermion hopping terms on any square lattice because U_h and U_v are diagonal operators in a basis for the Fock space obtained from a momentum basis for single-particle states).

The sequence length can be reduced by combining $U_U \left(\frac{\theta_U^{b+1}}{2} \right) U_U \left(\frac{\theta_U^b}{2} \right) = U_U \left(\frac{\theta_U^b + \theta_U^{b+1}}{2} \right)$; other orderings, such as first applying odd terms and then applying even terms will slightly reduce the gate depth.

It is important to emphasize that we are *not* concerned with Trotter error other than for a desire to preserve quantum numbers such as momentum; while this choice of unitaries gives evolution similar to the evolution under $H(b)$, we are optimizing the parameters for evolution with these unitaries. If we instead did exact evolution under a time-varying Hamiltonians (which is possible for atomic quantum gases in optical lattices [13]), we would instead optimize to a different optimum in the parameters. We expect that the difference in these evolutions can be largely absorbed into a small shift in the variational parameters. Alternatively, one can think that instead of the terms h_{a_i} in Eq. (1) being chosen from the terms h_U, h_v, h_h they correspond to the terms h_U, h_v , as well as all individual hopping terms that sum up to h_h , with all those hopping terms using the same parameter in a given step.

D. Optimization with Exact Energies

We first consider the case where we can exactly calculate energies on a classical computer and use a simple optimization procedure. Six points (a “point” is a set of parameters) are chosen at random near the origin. For each point, we first follow a greedy noisy search, where we slightly perturb the values of the points, accepting whenever this reduces the energy, for a total of 150 evaluations of the energy. We then use Powell’s conjugate direction [14] method until it converges. After following this procedure for each of the six points, we keep the point whose energy is lowest at the end of the procedure, and we discard the other five points. For the point we keep, we alternate greedy noisy search and Powell search until neither can find an improvement. Our greedy noisy search uses a simple algorithm to determine step size: every thirty trials we count the number of acceptances. If that number is large, the step size is increased; if the number is small, the step size is reduced.

We call this optimization algorithm the “global variational” method as it involves optimizing all parameters simultaneously. Using the global variational approach and an ansatz with $S = 3$, we obtain energy errors of 2.0×10^{-8} , 0.019, 0.029, 0.083, and 0.59 for $N = 4, 6, 8, 10$, and 12. Increasing S improves the error, but at the larger sizes convergence was slow and results varied greatly from run to run, suggesting that the minimization was getting stuck in local minima. One clue to the origin of this difficulty is that in some cases the energy reduced but the overlap with the ground state also reduced. This is possible if it also reduces the amplitude of some highly excited state but increases the amplitude of a low excited state.

To overcome this problem, we used an alternative procedure to find the optimum. This procedure is inspired by annealing and we call it the “annealed variational” method. Let $H_s = th_H + th_V + sUh_U$ interpolate from H_0 to H_1 by changing the coupling U . We use the same ansatz (3), but we use a different procedure to select a starting point for the search. We first do a single step ansatz using as Ψ_I the ground state of H_0 and targeting $H_{1/S}$, using the same optimization as above, calling the resulting parameters θ_X^1 . We then use the Ψ_T from that optimization as Ψ_I for another single step targeting $H_{2/S}$, giving θ_X^2 . We continue, targeting $H_{3/S}, \dots$, using Ψ_T from one step as Ψ_I for the next. This then gives a trial state using S steps for H_1 using parameters θ_X^b for $b = 1, \dots, S$. We call this state the result of sequential optimization. We use those parameters as the starting point for a further search as before (i.e., a global variational search using the sequential optimization as the starting point), calling the resulting parameters the result of full optimization. (An alternate method that we tried that in some cases worked better was to do a similar anneal, but to target H_1 on every step; thus, one would parameters for a single step ansatz targeting H_1 ; then use that to do another step again targeting H_1 , and so

S	ΔE^s	ΔE^f	P^s	P^f	ΔE^s	ΔE^f	P^s	P^f
	4				6			
3	0.24	1.00×10^{-8}	0.7180	1.0000	0.062	0.033	0.9821	0.9903
5	0.20	3.00×10^{-8}	0.76289	1.0000	0.034	0.002	0.9912	0.9995
7	0.17	2.00×10^{-8}	0.8021	1.0000	0.018	0.00033	0.9954	0.9999
9	0.15	7.00×10^{-8}	0.8275	1.0000	0.013	0.00018	0.9967	1.0000
11	0.13	2.00×10^{-8}	0.8460	1.0000	0.012	0.00011	0.9970	1.0000
	8				10			
3	0.1	0.033	0.9790	0.9934	0.13	0.083	0.9217	0.9374
5	0.042	0.0046	0.9906	0.9983	0.087	0.041	0.9388	0.9585
7	0.031	0.0030	0.9930	0.9989	0.066	0.022	0.9492	0.9710
9	0.024	0.0013	0.9947	0.9995	0.053	0.014	0.9566	0.9809
11	0.019	0.00089	0.9960	0.9997	0.042	0.012	0.9626	0.9841
13	0.015	0.00038	0.9968	0.9999	0.036	0.0069	0.9660	0.9929
15	0.013	0.00031	0.9973	0.9999	0.033	0.0052	0.9676	0.9959
17	0.012	0.00022	0.9976	0.9999	0.032	0.0032	0.9682	0.9983
19	0.010	0.00027	0.9978	0.9999	0.032	0.0017	0.9688	0.9993
	8*				12			
3	0.66	0.53	0.3889	0.5231	0.69	0.59	0.3046	0.4102
5	0.56	0.17	0.4620	0.8727	0.59	0.39	0.3576	0.6154
7	0.49	0.065	0.5164	0.9353	0.54	0.18	0.3958	0.8520
9	0.44	0.046	0.5600	0.9501	0.51	0.087	0.4256	0.9294
11	0.40	0.032	0.5969	0.9609	0.48	0.054	0.4483	0.9541
13	0.36	0.022	0.6280	0.9685	0.46	0.035	0.4667	0.9714
15	0.33	0.017	0.6551	0.9829	0.45	0.025	0.4836	0.9790
17	0.31	0.010	0.6799	0.9910	0.43	0.021	0.4997	0.9823
19	0.28	0.0083	0.7030	0.9935	0.42	0.015	0.5152	0.9883

TABLE I: Performance using the annealed variational method. Numbers 4, 6, 8, 10, 12 indicate different values of N . 8* is described in text. The left-hand column is number of steps, S . Quantities $\Delta E^s, \Delta E^f$ indicate error in ground state energy after sequential and full optimization, respectively. Quantities P^s, P^f indicate absolute squared ground state overlap in those two cases, respectively.

on; this sometimes did better at a small number of steps and also did slightly better at the given number of steps on the $N = 12$ site model.)

The results of the annealed variational method are listed in Tab. I. For $S = 3, N > 4$, this algorithm is seen to be only marginally worse than the global variational method, but we found that it was faster at finding the optimum (the number of energy evaluations required depended on N , but the annealed variational method always required fewer than the original method, and in some cases required a factor of five times fewer evaluations). The most important advantage of the annealed variational method is that it becomes significantly more accurate than the global variational method for $S > 3$, where now the error drops consistently with increasing S , without the convergence issues seen using the global variational method. While generally larger N is more difficult, $N = 4$ shows worse performance than $N = 6, 8$, and 10 after sequential optimization and only shows better performance after full optimization, perhaps due to the degeneracy of the non-interacting Hamiltonian for $N = 4$.

An additional model shown in the table is called “8*”. This has eight sites with the same geometry as before, but we change the hopping strength in the horizontal direction to $1/\sqrt{2}$ and change the sign on the horizontal hoppings which are periodic, thus effectively inserting a π -flux for hopping around a loop. The resulting single particle spectrum has $N/2 - 2$ states below the Fermi energy and four states at the Fermi energy, giving the same many-body ground state degeneracy as for $N = 12$ sites. The flux model and $N = 12$ show very similar errors and are both distinctly more difficult than the others. Nonetheless, using only 33 parameters we obtain over 95% overlap and with 45 parameters 97.9% overlap for $N = 12$, despite the Hilbert space having a much larger dimension: $\binom{12}{6}^2 = 853776$ at the given filling. For $N = 10$, a 93.7% overlap is obtained with only 9 parameters, in a 63504-dimensional space (symmetries such as total spin slightly reduce this dimension).

E. Inexact Optimization: Gate Count and Run Time

Next we consider the effects of sampling error assuming one measures individual terms in the Hamiltonian on a quantum computer, considering a tradeoff between run-time and accuracy. Clearly, given a large enough number of samples, one can reduce the statistical error to the point that one can obtain the same accuracy as above. However, for a smaller number of samples, it is useful to modify the optimization algorithm. The trouble is that a small change in a point leads to only a small change in energy, and it thus requires a large number of samples to determine whether or not the energy improves.

We thus used a different algorithm and tested it for $N = 8$. Starting with all parameters equal to zero we randomly order the parameters and try increasing or decreasing each parameter by a constant, looking for an improvement. This is repeated (with parameters re-ordered again randomly) with slightly changed constants, until no further improvement is found. To determine if there is an improvement, we sample at the given points until the difference between the energies becomes twice the standard deviation (or until a maximum number of samples is done). On an average over ten runs, we were able to obtain more than 98% overlap with the ground state using 506 different point evaluations and $8.5 \cdot 10^5$ average samples per point, for 4.3×10^7 total samples.

Now we wish to estimate the number of gates required to obtain a single sample, including preparing the state Ψ_I , implementing the unitaries, and finally measurement. Ψ_I is a Slater determinant and can be prepared using Givens rotations [12] (using fewer gates than the strategies in Ref. 15). The unitaries U_U, U_h, U_v can be implemented efficiently using Jordan-Wigner cancellation techniques [16]. Finally, for measurement, the terms in the Hamiltonian can be broken into at most four commuting sets. For $N = 4n = 8, 12, \dots$ these are the Hub-

bard terms, the vertical hopping terms, and two sets for the horizontal hopping terms. Slightly more complicated divisions of the hopping terms are needed for $N = 4n + 2 = 6, 10, \dots$, while a model with more horizontal rows will require five sets.

The cost of implementing the unitaries dominates the measurement cost. We can measure h_U with a cost (in gate count) that is almost identical to the cost to implement U_U ; see, for example, the discussion around Fig. 12 of Ref. 12 and references therein. Similarly, we can measure h_v with cost almost identical to the cost to implement U_v , while we can measure both sets of commuting terms in h_h with cost almost identical to the cost to implement U_h . Hence, since in a single run we only measure *one* of these four sets of terms (i.e., h_U or h_v or one of the two terms in h_h , assuming $N = 4n$), the measurement cost in a single run is roughly one-quarter the cost of a single step $U_U(\frac{\theta_U^b}{2})U_h(\theta_h^b)U_v(\theta_v^b)U_U(\frac{\theta_U^b}{2})$. In general, the cost of implementing $\prod_{b=1}^S \left(U_U(\frac{\theta_U^b}{2})U_h(\theta_h^b)U_v(\theta_v^b)U_U(\frac{\theta_U^b}{2}) \right)$ scales linearly with S .

The Slater determinant can be prepared using Givens rotations and the fast Fourier strategy [12], requiring on the order of $N \log_2(N)$ Givens rotations. These rotations will involve sites which are further separated than the nearest neighbor sites which appear in U_h, U_v , requiring longer Jordan-Wigner strings; ignoring the cost of the strings, the cost to execute these rotations will be comparable to the cost to execute $\log_2(N)$ rotations by U_h or U_v . For fixed S , this would eventually dominate at large N , but we expect that S will need to increase with N too and that the dominant cost will remain the cost of implementing the steps. Consider the case of $N = 8$, for example. The initial ground state has two particles in the zero momentum state in the horizontal direction, one in the symmetric state in the vertical direction and one in the anti-symmetric state. There are also two particles in momentum states $\pm\pi/2$ in the horizontal direction, in the symmetric state in the vertical direction. We can prepare the ground state as follows. Label sites in the top row 1, 2, 3, 4 and label sites in the bottom row 5, 6, 7, 8 with 1 directly above 5, and so on. Initialize with particles in sites 1, 3, 4, 5. Then, applying Givens rotations between pairs of sites 3, 7; 4, 8 to place particles initially in sites 3, 4 in the symmetric state in the vertical direction. Next apply Givens rotations between pairs 1, 2; 5, 6 to bring the particle in site 1 into a symmetric state between 1, 2 and the particle in site 5 into a symmetric state between 5, 6. Again apply Givens rotations between pairs 1, 3; 2, 4; 5, 7; 6, 8 so that now the particle initially in site 1 is in a symmetric state between sites 1, 2, 3, 4 while the particle in site 5 is in a symmetric state between sites 5, 6, 7, 8. Thus, we have successfully occupied both states with zero momentum in the horizontal direction. This same procedure in fact also produces the desired particles in momentum states $\pm\pi/2$ so it prepares the desired ground state. This uses a total of 8 Givens

rotations. The gate count cost is comparable to that to implement U_h and U_v .

For the case of quantum chemistry discussed below, the initial state is much simpler. It is a product state and so the cost to prepare is negligible compared to that to implement the unitaries.

A gate count estimate shows that about 1000 gates are required for $S = 2$ at $N = 8$ for a single run. This includes the cost to prepare Ψ_I , to implement the unitary rotations to create Ψ_T , and to perform measurement. To measure all terms, this must be multiplied by 4 as each run will give a measurement of the terms in *one* of the four commuting sets of terms. To give a rough estimate, if we ignore the possibility of parallelizing the circuits and assume a gate time of $1\mu\text{s}$ the total time would be $4.3 \times 10^7 \times 4 \times 1000 \times 10^{-6}$ seconds, or roughly 47 hours. This could be further improved using several devices in parallel to sample. Note that any given run requires only about 1000 gates, and thus poses only moderate demands on gate fidelities.

III. QUANTUM CHEMISTRY

A. The electronic structure Hamiltonian

We finally apply the method to three problems in quantum chemistry, where there are more terms than in the Hubbard model. Using the Born-Oppenheimer approximation by assuming that the nuclei behave classically and are localized, the Hamiltonian for the electronic degrees in second quantized form

$$H = \sum_{pq} h_{pq} c_p^\dagger c_q + \sum_{pqrs} h_{pqrs} c_p^\dagger c_q^\dagger c_r c_s \quad (4)$$

has the most general form for two-body interactions, due to the long range nature of the Coulomb interaction. The indices p, q, r and s refer to spin-orbitals, combining the orbital and spin index.

As most basis sets used in quantum chemistry are non-orthogonal we follow the standard procedure [17] of first performing a (classical) Hartree-Fock calculation and then use the Hartree-Fock orbitals as an orthogonal basis set. On classical computers we can only simulate quantum algorithms for very small basis molecules.

Specifically we consider HeH^+ which has *two* electrons, using a P321 basis with $N_{\text{SO}} = 8$ spin orbitals, H_2O which has 10 electrons, in an STO-6G basis with $N_{\text{SO}} = 14$, and BeH_2 which has 6 electrons in a basis with $N_{\text{SO}} = 14$. We use the Psi4 [18] quantum chemistry package to perform the Hartree Fock calculation and compute the matrix elements of the Hamiltonian (4). We also consider artificial hydrogen chains H_N , where we space hydrogen atoms along a line with distances of either 0.4614\AA or 2\AA . We used the global variational approach to optimize parameters for HeH^+ , H_2O , and BeH_2 and we used the annealed variational approach to optimize

S	$ \Delta E$ [mHa]	P	$ \Delta E$ [mHa]	P	$ \Delta E$ [mHa]	P
	HeH^+		H_2O		BeH_2	
1	7.8	0.9970	23	0.9886	22	0.9825
2	1.7	0.9992	9.5	0.9950	6.6	0.9935
3	0.59	0.9998	7.6	0.9955	6.3	0.9937
4	0.26	0.9999	6.8	0.9959	5.8	0.9939
5	0.088	0.9999	3.2	0.9980	4.23	0.9954
6	0.14	0.9999	3.1	0.9982	1.85	0.9977

TABLE II: Variational Method applied to HeH^+ and H_2O . ΔE is the error in the ground state energy, P is absolute squared overlap with ground state. Note that the energy is slightly worse for HeH^+ at $S = 6$ than at $S = 5$; this indicates that the optimization algorithm did not find the true minimum in this case.

H_N . All numbers reported for all methods are the best result obtained from three runs; due to randomness in the search the results differ slightly between runs.

B. Variational ansatz for quantum chemistry

We considered two different ansatzes for quantum chemistry. In the first ansatz, used for HeH^+ , H_2O , BeH_2 , instead of using $(O)(N_{\text{SO}}^4)$ variational parameters in our ansatz, we group the matrix elements into three terms $H = H_{\text{diag}} + H_{\text{hop}} + H_{\text{ex}}$, where

$$H_{\text{diag}} = \sum_p \epsilon_p c_p^\dagger c_p + \sum_{p,q} h_{pqqp} c_p^\dagger c_p c_q^\dagger c_q \quad (5)$$

is a sum of diagonal terms,

$$H_{\text{hop}} = \sum_{pq} h_{pq} c_p^\dagger c_q + \sum_{pq} h_{prrq} c_p^\dagger c_q c_r^\dagger c_r \quad (6)$$

contains normal and correlated hopping terms and

$$H_{\text{ex}} = \sum_{pqrs} h_{pqrs} c_p^\dagger c_q^\dagger c_r c_s \quad (7)$$

for p, q, r, s all distinct contains all other exchange terms.

As our final ansatz we use

$$\Psi_T = \prod_{b=1}^S \left(U_{\text{ex}} \left(\frac{\theta_{\text{ex}}^b}{2} \right) U_{\text{hop}} \left(\frac{\theta_{\text{hop}}^b}{2} \right) \times U_{\text{diag}} \left(\theta_{\text{diag}}^b \right) U_{\text{hop}} \left(\frac{\theta_{\text{hop}}^b}{2} \right) U_{\text{ex}} \left(\frac{\theta_{\text{ex}}^b}{2} \right) \right) \Psi_I, \quad (8)$$

where $U_X(\theta)$ approximates $\exp(i\theta H_X)$ for $X \in \{\text{ex, hop, diag}\}$, up to Trotterization error. The initial state Ψ_I is chosen to be the ground state of H_{diag} and the basis is a Hartree-Fock basis so that $H_{\text{hop}} \Psi_I = 0$. The unitary $U_{\text{diag}}(\theta) = \exp(i\theta H_{\text{diag}})$ can be implemented exactly since all terms commute. We implemented the unitary $U_{\text{hop}}(\theta)$ using an ‘‘interleaved’’ term ordering [16],

approximating it as a product $U_{\text{hop}}(\theta) = \prod_{p<q} U_{pq}(\theta)$ in arbitrary order, with

$$U_{pq}(\theta) = \exp \left[i\theta \left(h_{pq} c_p^\dagger c_q + \sum_r h_{prrq} c_p^\dagger c_r^\dagger c_r c_q + h.c. \right) \right], \quad (9)$$

where all terms in the exponential commute with each other. The term U_{hop} appears twice in Eq. (8) for each step; in a slight abuse of notation, the order of terms was reversed in the two different applications. The unitary $U_{\text{ex}}(\theta)$ was approximated as a product over factors

$$U_{pqrs}(\theta) = \exp \left[i\theta (h_{pqrs} c_p^\dagger c_r^\dagger c_r c_s + h.c.) \right], \quad (10)$$

in an arbitrary order. The term U_{ex} also appears twice in Eq. (8) for each step; in another slight abuse of notation, the order of terms was reversed in the two different applications. Thus, in each step the unitary applied is a second-order Trotter approximation to evolution under $H(b) \equiv \theta_{\text{ex}}^b H_{\text{ex}} + \theta_{\text{hop}}^b H_{\text{hop}} + \theta_{\text{diag}}^b H_{\text{diag}}$.

The second ansatz was used for the H_N chains. In this case we used four parameters. We did this by further dividing H_{ex} into a sum of two terms $H_{\text{ex}} = H_{\text{o}} + H_{\text{rest}}$, where H_{o} includes the terms in H_{ex} which do not annihilate the Hartree-Fock state (the “o” indicates that this is based on occupancy of the orbital in the Hartree-Fock state) and H_{rest} includes the remaining terms. We then used separate parameters to control $H_{\text{o}}, H_{\text{rest}}$. Thus, the final ansatz was

$$\Psi_T = \prod_{b=1}^S \left(U_{\text{o}} \left(\frac{\theta_{\text{o}}^b}{2} \right) U_{\text{rest}} \left(\frac{\theta_{\text{rest}}^b}{2} \right) U_{\text{hop}} \left(\frac{\theta_{\text{hop}}^b}{2} \right) \right. \\ \left. \times U_{\text{diag}} \left(\theta_{\text{diag}}^b \right) U_{\text{hop}} \left(\frac{\theta_{\text{hop}}^b}{2} \right) U_{\text{rest}} \left(\frac{\theta_{\text{rest}}^b}{2} \right) U_{\text{o}} \left(\frac{\theta_{\text{o}}^b}{2} \right) \right) \Psi_I, \quad (11)$$

Our results for the three benchmark molecules using various numbers of steps S are shown in Table II.

C. Unitary Coupled Cluster Ansatz

We have also performed simulations using UCC [9] for comparison as well as variants. UCC method fixes $\Psi_T = \exp(T)\Psi_I$, for T an anti-Hermitian operator. The variational parameters are contained in the choice of T . Ref. 8 proposes using a Trotter-Suzuki method to implement the transformation $\exp(T)$. This increases the depth of the circuit, which we wish to avoid. We have found in our simulations that more accurate results are in fact obtained using large Trotter steps, with the most accurate results obtained using between two and four second-order Trotter-Suzuki steps; all results reported below are for two such steps. The method thus differs from the usual UCC, although we continue to refer to it as such. We used similar optimization procedures as before, exactly evaluating (to numerical error) the energy as a function of parameters.

We set

$$T = \sum_{p<q} (T_{pq} c_p^\dagger c_q - h.c.) + \sum_{p<q,r<s} (T_{pqrs} c_p^\dagger c_q^\dagger c_r c_s - h.c.), \quad (12)$$

using all possible quadratic and quartic terms which are compatible with the symmetries of the system. We considered two cases, taking the parameters T_{pq}, T_{pqrs} either all real or all imaginary. More accurate results were found using real choices; this is likely due to the fact that the wavefunction is real. More flexibility could be obtained by using general complex values, but at the cost of a further increase in parameters. In the real case, we can drop diagonal terms such as $c_p^\dagger c_p$. Thus, all runs reported are for the choice of real parameters.

Further, in the UCC method, one considers only terms in which all creation operators act on orbitals which are unoccupied in the Hartree-Fock state and all annihilation operators act on orbitals which are occupied, or vice-versa as required by Hermiticity; equivalently, these are the terms which do not annihilate the Hartree-Fock state. In our numerics, we considered three alternative possibilities for a total of four possible choices of terms to keep. One choice we call RAA, in which *all* terms are kept (the “A”s refer to keeping all quadratic and all quartic terms, while the “R” refers to T being real). Another choice we call ROO, in which now we keep only quadratic and quartic terms which do not annihilate the Hartree-Fock state (the “O” refers to “occupation”, as whether or not the term is kept depends upon the occupation number). The choice ROO is precisely the same as UCC. The next choice we term RNO, where *no* quadratic terms are kept and quartic terms are kept only if they do not annihilate the Hartree-Fock state (the “N” refers to “none”); this choice has the fewest terms. The last choice is RAO, where all quadratic terms are kept and only quartic terms are kept if they do not annihilate the Hartree-Fock state. Thus, in decreasing order of number of terms kept, they are RAA, RAO, ROO, RNO. The variants can be collectively termed Rxx.

For molecules, to find terms compatible with symmetries, we included all terms which had nonzero coefficients in the Hamiltonian (this should work for small molecules, while for large molecules some terms compatible with symmetry will have a small enough coefficient that they get truncated). The results are shown in Table ?? . Note that for $\text{HeH}^+, \text{H}_2\text{O}$, all methods are able to find the ground state to very high accuracy. This may be due to the large number of parameters compared to Hilbert space dimension. For RAA, for HeH^+ , there are 192 parameters in the imaginary case (almost as many in the real case) compared to a Hilbert space dimension of 64 with the given number of up and down electrons (possibly further reduced by symmetry) while for H_2O there are 595 parameters compared to a Hilbert space dimension of 441. This is an obstacle that we encounter when trying to simulate small molecules on a classical computer. The number of variational parameters in this method grows as N_{SO}^4 while the Hilbert space dimension grows exponen-

	HeH ⁺	H ₂ O	BeH ₂
RAA	3.2×10^{-4}	3.3×10^{-2}	0.16
RAO	5.3×10^{-3}	4.4×10^{-2}	0.41
ROO	1.3×10^{-2}	5.9×10^{-2}	0.42
RNO	0.42	0.36	0.66

TABLE III: Rxx methods applied to various small molecules. Table shows energy error in mHa. Note that ROO is the same as UCC.

tially with N_{SO} at fixed filling fraction; eventually the exponential growth beats the polynomial (indeed, this is the whole reason for interest in a quantum computer), but since the number of parameters grows at such a high power of N_{SO} , one needs fairly large systems to see this. The other methods have smaller numbers of parameters; some slight accuracy loss can be seen in RNO.

The reduction in the number of terms when going from RAA to RNO or ROO depends upon the specific molecule. At half-filling, for molecules with N_{SO} spin-orbitals with $N_{SO} \gg 1$, the number of quartic in RAA is proportional to N_{SO}^4 and will be much larger than the number of quadratic terms. Roughly 1/8-th of these terms will be retained in RAO, RNO, or ROO; to see this, note that a term is retained if p, q both correspond to unoccupied orbitals and r, s both to occupied (which occurs with probability $1/2^4 = 1/16$ if one selects p, q, r, s randomly at half-filling) or p, q are both occupied and r, s are both unoccupied (which also occurs with probability $1/16$). This leads thus to a constant factor gain; away from half-filling the constant factor improvement becomes larger.

For BeH₂, the Hilbert space dimension is 1225, compared to 595 terms for RAA and so the Hilbert state dimension is smaller than the number of terms, but still comparable. For BeH₂, RAA is able to achieve an energy error of 0.157 mHa, at the cost of over 2×10^5 energy evaluations. A smaller number of evaluations leads to reduced accuracy; the accuracy improves most rapidly up to roughly 2×10^4 evaluations with an error of roughly 0.5 mHa at that point. In contrast, the Hamiltonian variational method used between 5000 – 10000 evaluations for the data above. For a given number of evaluations in the optimization procedure, in general RAA had slightly improved energy than the methods with fewer terms.

D. Hydrogen chains

Real applications will be concerned with molecules with larger N_{SO} . In order to access this regime we turn to a different system, H_N chains, varying the number of atoms from $N = 2$ to $N = 10$; this allows us to test scaling by considering a sequence of different system sizes, all at fixed filling fraction and with similar couplings. For $N = 2, 4$ and 6 the number of terms in RAA exceeds the Hilbert space dimension, while for $N = 8$ there are

	ΔE [mHa]	P	ΔE [mHa]	P	ΔE [mHa]	P
	H ₆		H ₈		H ₁₀	
RAA	0.4	0.995	51	0.786	152	0.578
RAO	8.03	0.976	54	0.797	89	0.674
ROO	10	0.979	54	0.787	80	0.697
RNO	13	0.974	50	0.817	69	0.746
H Var 3	21	0.947	35	0.904	51	0.892
H Var 6	3.0	0.982	7.3	0.959	11	0.960
H Var 9	0.97	0.984	2.6	0.965	3.6	0.967

TABLE IV: Variational methods applied to hydrogen chains. Left-hand column gives method. Rxx refers to Rxx methods with ROO=UCC while H Var 3, 6, 9 refer to Hamiltonian variational with $S = 3, 6, 9$.

2964 terms with a Hilbert space dimension of 4900 and for $N = 10$ there are 7230 terms with a Hilbert space dimension of $N = 63504$. We performed simulations at separations of either 0.4614\AA or 2\AA between the atoms. For the smaller spacings, there is a much higher overlap between the true ground state and the Hartree-Fock ground state, while for the larger spacing, the overlap becomes much smaller; this is similar to the Hubbard chain at small U/t compared to large U/t .

The results are shown in Table IV. In general, for larger N , the Hamiltonian variational method is able to obtain significantly higher overlap (the use of four parameters improves this; with only three parameters, a larger S is required to attain the same accuracy).

For H₆, increasing the number of terms in Rxx leads to improved performance: RAA outperforms RAO which outperforms ROO which outperforms RNO on measures of energy. However, for H₈, H₁₀, the reverse is true, with RNO having the highest overlap of any Rxx method. When we consider the energy accuracy as a function of number of evaluations, we find that at a given number of evaluations RAA generally performs better (often by only a small amount). However, the methods with fewer parameters are often able to continue the optimization out for more steps before getting stuck. We do not fully understand the mechanism for this, but one possibility is that once the optimization routine gets very close to a local minimum, it is likely that motion in a random direction will hurt the energy. Methods with more parameters are more likely to move in an incorrect direction simply due to the increased dimensionality of the search space and hence will be more likely to fail to find an improvement once very close to the minimum. One possible improvement then would be to run RNO until it found its optimum; then, take that endpoint as a starting point for an ROO or RAA search; this requires investigation.

An interesting further feature of the convergence is that the improvement in energy is rather rapid initially and then transitions to a routine with much slower improvement. For example, for H₁₀, the energy for the various Rxx versions has an error of ~ 0.2 after roughly 5×10^4 evaluations, followed by a very gradual improve-

ment, with the RNO numbers above occurring after over 2.5×10^5 evaluations.

As N increases, we find that the annealed variational approach gets increasingly better than the UCC; the UCC shows good performance at small sizes, perhaps due to the large number of parameters compared to the Hilbert space dimension, but gets increasingly less accurate at large sizes and is more difficult to optimize. For the smaller spacing, Rxx performed relatively better but a similar trend was found.

A more flexible method for quantum chemistry is to use multiple steps (as in the Hamiltonian variational approach) but allow every term to have a separate parameters (as in UCC); we call such method TRxx for “time-dependent Rxx”, where the “time” refers to different steps. Optimizing all these parameters may become difficult. Note that for Hubbard, the unitary coupled cluster method requires a large increase in circuit depth compared to the Hamiltonian variational approach due to all the additional terms in T and hence it is likely to not be suitable for that context unless particularly useful choices of a small number of terms in T can be found. The Hamiltonian variational approach can be regarded as a useful choice of which terms to take and how to vary them.

IV. RESOURCE ESTIMATES FOR PRACTICAL APPLICATIONS

We want to end with a discussion about resource requirements for practical applications that might go beyond what can currently be done classically.

A. Hubbard model

For the Hubbard model interesting results can be obtained for lattice sizes beyond 10×10 , and thus with $N \geq 100$ sites, which requires about 200 qubits (or slightly more for ancillas if we want to parallelize the circuits). As the circuits to implement the various terms in the Hamiltonian can be efficiently parallelized [12], the parallel circuit depth will not substantially increase except for the potential need to go to a larger number of steps S . In order to be competitive to the best classical variational wave functions and distinguish competing ground states (such as stripe order versus uniform superconducting ground states) we need to aim for an error of $\epsilon \approx 10^{-3}t$ in the energy *per site* [19].

Assuming a variance of order 1 for the measurement of the hopping terms $c_{i,\sigma}^\dagger c_{j,\sigma} + c_{j,\sigma}^\dagger c_{i,\sigma}$ and a reduced variance of order $1/U$ for the double occupancy term $c_{i,\uparrow}^\dagger c_{i,\uparrow} c_{i,\downarrow}^\dagger c_{i,\downarrow}$ (due to a suppression of double occupancy to about t/U for large U) we get an error estimate of

$$\epsilon^2 \approx U^2 \frac{t/U}{MN} + 4t^2 \frac{1}{MN}, \quad (13)$$

assuming M measurements for each of five terms (double occupancy, horizontal and vertical hopping for each of the spins) and using that we can perform N measurements in parallel in each of the runs. This is consistent with the number of samples that was found to be necessary comparing this to the values measured for small systems in Sec. II E. Using relevant values of $t = 1$, $U = 8$, $N = 100$, we obtain $M \approx 12/\epsilon^2 N \approx 120,000$ samples for each of the five terms, or about a total of 600,000 samples per energy evaluation. Again assuming gate times of order $1\mu s$, as in Sec. II E, the total estimated times needed to achieve convergence remains of the order of days when parallelizing the circuits, even considering that larger S and more energy evaluations might be required. While not trivial on a small quantum computer, this is potentially practical in the not too distant future.

B. Quantum chemistry

For quantum chemistry applications the required resources will be more demanding. Writing the Hamiltonian as $H = \sum_{i=1}^{N_{\text{terms}}} h_i \mathcal{O}_i$, where N_{terms} is the number of terms, and the numbers h_i the coefficients of the terms \mathcal{O}_i we obtain for the error

$$\epsilon^2 = \sum_i \frac{|h_i|^2 \text{Var}(\mathcal{O}_i)}{M_i}, \quad (14)$$

assuming M_i measurements of each of the operators \mathcal{O}_i , where $\text{Var}(\dots)$ denotes the variance in the measurement of the operator for the given trial state. Minimizing the error for a total number of measurements M , we choose $M_i \propto |h_i|$, and bounding the variances by $\text{Var}(\mathcal{O}_i) \leq 1$ we get

$$M \approx \frac{(\sum_i |h_i|)^2}{\epsilon^2} \quad (15)$$

As the variances of the (large) diagonal terms ϵ_p and h_{pqpp} will be small for orbitals where the occupation number is close to 1 or to 0 (this holds for all orbitals in the molecules studied above), we will drop these terms from the error estimates and consider just the off-diagonal terms in the following order of magnitude estimate. Some of the off-diagonal terms have small variance. For example, any term h_{pq} in which p, q both have occupation number close to 1 or close to 0 will have small variance. However many off-diagonal terms have large variance; for example, a term h_{pq} where p has occupation number close to 1 and q has occupation number close to 0 has large variance. Hence, as a rough estimate, we treat the variance of the off-diagonal as a constant of order unity.

We find for the sums of matrix elements $\sum_i |h_i|$ the values 11.3Ha for HeH^+ , 12.3Ha for BeH_2 and 36Ha for H_2O , which results in 10^8 to 10^9 required samples for each energy evaluation to achieve an error of 1mHa. This is about a factor 1000 larger than for the Hubbard model, which already needed on the order of a few days to be

optimized at an assumed $1\mu\text{s}$ gate time and measurement time. Considering additionally that the circuits are more complex, this will require a massively parallel cluster of quantum computers to perform the simulations for small molecules in a reasonable time.

Moving from these toy problems that can easily be simulated classically to a more challenging problem, such as Fe_2S_2 in a small single particle basis (STO-3G with $N = 112$ spin-orbitals) we have $\sum_i |h_i| \approx 3\text{kHa}$, thus requiring about 10^{13} samples per energy evaluation. Even assuming that still only about 10^6 energy evaluations could be sufficient to optimize the ansatz, the required number of 10^{19} total samples is too large for the variational algorithm to be practical in this form. For this model, a single instance of the Hamiltonian circuit is on the order of 2×10^8 gate executions, which is an estimate of the cost to prepare a sample using either a UCC method truncated at all four fermi terms or the Hamiltonian variational method. This amounts to a total of 10^{26} gate operations assuming 20 – 30 commuting terms are measured on each sample. This estimate assumed that all of the energy evaluations in the optimization were conducted at the same accuracy as the final estimate; potentially some of the earlier energy evaluations could be conducted at lower energy accuracy, but following a noisy search to the minimum may require that many steps of the evaluation be conducted at *higher* accuracy than the final evaluation, since we have found that often only a very small improvement in energy is obtained on each step of the search.

V. DISCUSSION

We have presented a method of constructing variational states for arbitrary Hamiltonians, and tested it numerically. Preparation of quantum states using short quantum circuits is likely the lowest hanging fruit for small quantum computers to achieve quantum supremacy, and outperforming classical computers. No expensive quantum error correction may be needed if the gate fidelity is high enough to allow a short quantum circuit to succeed in preparing a state.

We find that a modest number of parameters enables us to obtain a very large overlap with the ground state, helping reduce the number of evaluations required to obtain good values of the parameters. Of course, the minimum possible circuit depth is attained simply by declaring the class of variational states to be “all states that can be constructed with unitary quantum circuits of given depth”. However, this would require an impractical optimization. The method here allows us a flexible class of states with a small number of parameters, with the circuits chosen from the terms in the Hamiltonian. We find that for larger systems with stronger interactions, this approach is able to obtain significantly higher overlap than ansatz wave functions based on the unitary coupled cluster formalism or its Rxx variants.

Although the state preparation is fast, the number of measurements required to estimate the energy with sufficient accuracy is a huge challenge. While the estimates for interesting quantum chemistry applications are astronomical and the variational algorithm for these problems thus impractical in its current form, the demands for the Hubbard model are less challenging. There, the smaller number of terms, limited energy range and translation invariance help reduce the required number of measurements to make a quantum variational approach potentially practical, although still demanding.

We have also given a detailed numerical simulation of UCC and Rxx for a variety of systems, obtaining better understanding of its convergence properties. Further, we have shown that small Trotter numbers suffice, reducing circuit depth for this algorithm.

A less demanding application may be to use these ansatz wave functions for state preparation, which uses shorter circuits than would be needed if one instead prepared the state by adiabatic preparation. We have attempted to compare this by optimizing the parameters in an anneal for the Hubbard model: we consider an anneal using a linear path from the model at $U = 0$ to the final model at $U = 2$. This linear path was implemented by discrete time steps dt on the quantum computer, evolving using a second order Trotter-Suzuki for each time step (hence, the individual rotations in this Trotter-Suzuki evolve by time step $dt/2$). Thus, there are two parameters that quantify the anneal: the total annealing time T , and the time step dt . The ratio T/dt is the number of second order Trotter-Suzuki steps that must be implemented. We optimized these two parameters separately to obtain the minimum ratio that achieved 0.99 or higher overlap with the ground state. For the $N = 8$ model, this minimum was 8, to be compared to $S = 3$ to attain this accuracy using the variational method. For $N = 12$, the minimum was 640, to be compared to $S = 19$ using the variational method (the table shows only 0.9883 overlap for $S = 19$, but slightly higher overlap was obtained by targetting H_1 on every step in this case). We expect that larger system sizes may lead to even more significant gains in depth.

This state preparation may be useful for measuring properties of the states, in particular when combined with the quadratic speedup of Ref. 12. Additionally, this state preparation may be useful for annealing to even larger states. For example, if we construct a short circuit to create a Hubbard model ground state on 12 sites (and indeed, one outcome of this work is that we have been able to find such a circuit using a classical computer), then one could create two or more copies of this ground state and then anneal from the Hubbard Hamiltonian describing 24 sites partitioned into 2 decoupled sets of 12 sites to the Hubbard Hamiltonian for 24 coupled sites. This is an example of the “joining” procedure of Ref. 12, using the variational method to simplify the creation of the decoupled systems.

We optimized the parameters by a search procedure

which treated the energy as a function of parameters of a black-box function to be optimized. It is in fact possible to also determine the derivative of this function using a quantum circuit, but this requires roughly doubling the depth of the circuit. It is not clear whether or not access to the derivative would improve the optimization.

One might further speculate whether optimizing over parameters by searching would be useful for the optimization problems of Refs. 10, 11, improving over the performance of an adiabatic algorithm. Further numerical work may help answer this.

Acknowledgments

We thank P. Love, J. McClean, and A. Aspuru-Guzik for discussion and useful explanations of UCC as in Refs. 8, 9. This work was supported by Microsoft Research. MH and MT acknowledge hospitality of the Aspen Center for Physics, supported by NSF grant PHY-1066293.

-
- [1] U. Schollwoeck, *Ann. Phys.* **326**, 96 (2011).
 - [2] G. Vidal, *Phys. Rev. Lett.* **101**, 110501 (2008).
 - [3] F. Verstraete and J. I. Cirac, *cond-mat/0407066* (2004).
 - [4] M. Schwarz, K. Temme, and F. Verstraete, *Phys. Rev. Lett.* **108**, 110502 (2012).
 - [5] M. Schwarz, T. S. Cubitt, K. Temme, F. Verstraete, and D. Perez-Garcia, *Physical Review A* **88**, 032321 (2013).
 - [6] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **98**, 140506 (2007).
 - [7] C. M. Dawson, J. Eisert, and T. J. Osborne, *Phys. Rev. Lett.* **100**, 130501 (2008).
 - [8] A. Peruzzo et. al., *Nature Comm.* **5**, 4213 (2014).
 - [9] A. Taube and R. J. Bartlett, *Int. J. Quant. Chem.* **106**, 33933401 (2006).
 - [10] E. Farhi, J. Goldstone, and S. Gutmann, *arXiv:1411.4028*.
 - [11] E. Farhi, J. Goldstone, and S. Gutmann, *arXiv:1412.6062*.
 - [12] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, *arXiv:1506.05135*.
 - [13] D. Jaksch and P. Zoller, *Annals of Physics*, **315** 52 (2005).
 - [14] M. J. D. Powell, *Computer Journal* **7**, 155 (1964).
 - [15] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, *Phys. Rev. A* **65**, 042323 (2002).
 - [16] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, *QIC* **15**, 1 (2015).
 - [17] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, *Molecular Physics* **109**, 735 (2011).
 - [18] “Psi4: An open-source ab initio electronic structure program,” J. M. Turney, A. C. Simmonett, R. M. Parrish, E. G. Hohenstein, F. Evangelista, J. T. Fermann, B. J. Mintz, L. A. Burns, J. J. Wilke, M. L. Abrams, N. J. Russ, M. L. Leininger, C. L. Janssen, E. T. Seidl, W. D. Allen, H. F. Schaefer, R. A. King, E. F. Valeev, C. D. Sherrill, and T. D. Crawford, *WIREs Comput. Mol. Sci.* **2**, 556 (2012). (doi: 10.1002/wcms.93).
 - [19] P. Corboz, T. M. Rice, and M. Troyer, *Phys. Rev. Lett.* **113**, 046402 (2014).