

Gate count estimates for performing quantum chemistry on small quantum computers

Dave Wecker,¹ Bela Bauer,² Bryan K. Clark,^{2,3} Matthew B. Hastings,^{2,1} and Matthias Troyer⁴

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

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

³*Kavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA*

⁴*Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland*

As quantum computing technology improves and quantum computers with a small but non-trivial number of $N \geq 100$ qubits appear feasible in the near future the question of possible applications of small quantum computers gains importance. One frequently mentioned application is Feynman's original proposal of simulating quantum systems, and in particular the electronic structure of molecules and materials. In this paper, we analyze the computational requirements for one of the standard algorithms to perform quantum chemistry on a quantum computer. We focus on the quantum resources required to find the ground state of a molecule twice as large as what current classical computers can solve exactly. We find that while such a problem requires about a ten-fold increase in the number of qubits over current technology, the required increase in the number of gates that can be coherently executed is many orders of magnitude larger. This suggests that for quantum computation to become useful for quantum chemistry problems, drastic algorithmic improvements will be needed.

I. INTRODUCTION

The excitement over quantum computation stems from the promise that quantum computers will be able to solve problems for which classical computers don't have enough resources. The evidence for this comes from the discovery of quantum algorithms [1–7] which, at least asymptotically, are exponentially faster than classical algorithms. This assures us that eventually, when sufficiently large quantum computers exist, they will fulfill this promise. On the flip side, simple quantum algorithms have already been performed: for example, the number 15 has been factored [8] and the energy of an extremely simple molecule has been calculated [9, 10]. Although an important first step, these quantum calculations are still deep in the regime accessible to classical computers. It is interesting, then, to explore what minimal resources are needed for quantum computers to solve problems that classical computers are unable to solve. In particular, this encourages us to explore problem instances which are just big enough to be outside the range of classical computers (say, for the next decade) and understand the quantum resources needed to solve these problems. We call these classically-intractable problems. In this work, we take up this task for the area of quantum chemistry.

Feynman's original proposal for a quantum computer [11] was motivated by the exponential complexity of simulating many classes of quantum systems on a classical computer. The wave function of N 2-level systems, (*e.g.* N spin-1/2 variables pointing up or down in a quantum magnet or N spin-orbitals in a molecule each being occupied with either 0 or 1 electrons) lives in the Hilbert space \mathbb{C}^{2^N} and thus needs an exponentially large number of 2^N classical variables to store. In contrast, on a quantum computer storing the same wave function requires

only N qubits. This reduces the memory requirement from exponential to linear and the runtime cost for many computations on the quantum system from exponential to polynomial.

The current state of the art in numerically exact classical algorithms, based on the diagonalization of the Hamiltonian matrix using standard linear algebra methods either in the full Hilbert space or in a large Krylov subspace, can reach approximately $N = 50$ spin orbitals [12–15]. Approximate methods for fermionic computation are starting to reach chemical accuracy on strongly correlated systems for up to $N = 70$ spin orbitals [16]. Hence, an interesting application of a quantum computer needs to reach at least $N = 50$ spin orbitals to offer any advantages over classical machines and realistically needs approximately $N = 100$ spin orbitals to be significantly more useful than current classical algorithms. To store a wave-function of this size requires full coherent control over at least 100 qubits. Experimentally, such systems seem feasible in the near-term future: Ion trap experiments have already demonstrated coherence and entanglement between fourteen qubits [17] and many more ions have been trapped, but not yet entangled. Using superconducting qubit technology, around 10 qubits can be controlled and a few hundred gates can be executed coherently.

While there has been great progress towards non-trivial quantum computers with a small number of qubits, the development of quantum algorithms and exploration of applications for such devices has lagged behind. Factoring large integers using Shor's algorithm [1] is the canonical application for quantum computers, but it requires many thousands of qubits to factor a number that cannot be factored by classical algorithms [18–20]. The electronic structure problem for molecules, on the other hand, seems a more natural place where non-trivial applications may exist for machines with a limited

number of qubits and a significant amount of literature has been devoted to this topic [5, 21–38]; for recent reviews, see Refs. [28, 39]. Like Shor’s algorithm, solving the electronic structure problem may also admit an exponential speedup but affords interesting possibilities with fewer necessary qubits. In addition, the technological benefits of quantum chemistry simulations are rich: For example, finding better catalysts to be used in many industrial-scale chemical processes – even including very basic processes, such as nitrogen fixation – has challenged researchers for decades [40]. In these problems, approximate approaches such as density-functional theory do not yield sufficient accuracy for the correlation energies, while more accurate methods, such as the density matrix renormalization group, have so far not been able to simulate sufficiently large systems in a reasonable time-frame.

While there are many important quantities that characterize molecular systems, in this paper, we focus particularly on the measurement of the electronic ground state energy. Computing these energies (and their respective derivatives) is a basic starting point for computing other observables. We implement one of the standard approaches to performing such calculations on a quantum computer [22, 25], for brevity henceforth referred to as *quantum full configuration interaction* (QFCI), at the level of the individual circuit elements and compare the results obtained for a water molecule (10 electrons, 14 spin-orbitals in an STO-3G basis) to those obtained by the equivalent standard full configuration interaction (FCI) calculation, validating the approach. We note that water in this basis is a standard example that was already considered in Ref. 22. We are then able to measure the costs involved in the quantum computation, including the number of qubits, number of circuit elements, and the circuit depth needed to perform this computation.

By systematically analyzing the effects of time-step errors and gate count, we show the scaling of the algorithm with the number N of spin orbitals to be $\mathcal{O}(N^9)$ if all gates are executed in a serial fashion, and $\mathcal{O}(N^8)$ if we allow for parallel execution of gates. Using the example of a water molecule, we also determine the prefactors involved and thereby set a baseline of requirements for a quantum computer to perform calculations on larger molecules.

II. THE COULOMB HAMILTONIAN IN QUANTUM CHEMISTRY

Using a Born-Oppenheimer approximation to fix the positions of the nuclei in space, the electronic structure problem for a molecule reduces to finding the low-lying spectrum of the electronic degrees of freedom. For a full configuration interaction approach, a basis of single particle orbitals, such as the STO-3G basis used here, must be chosen. The choice of basis here dictates the number of orbitals that need to be considered in the FCI calculation. One then obtains a basis of molecular orbitals by

performing a Hartree-Fock calculation and rewrites the Hamiltonian in a second-quantized form in terms of these orbitals, where it takes the form

$$H = \sum_{pq} t_{pq} c_p^\dagger c_q + \frac{1}{2} \sum_{pqrs} V_{pqrs} c_p^\dagger c_q^\dagger c_r c_s. \quad (1)$$

Here, c_p and c_p^\dagger denote the annihilation and creation operators for an electron in a set of $p = 1, \dots, N$ spin orbitals. For this paper we want to focus on calculating the ground state energy E_0 and generating the ground state wave function $|\psi_0\rangle$ of this Hamiltonian.

The energy scales for the molecular problems are set by the core energy of the atom, giving energies of approximately 100-1000 Hartree (E_h) for small molecules. In order to perform useful quantum chemistry, a chemical accuracy of approximately 1 milli-Hartree (mE_h) is important. This means we have to resolve the energy scales to one part in a million.

Currently, quantum chemists use the exact full configuration interaction (FCI) method, and a variety of approximate methods to tackle this problem. Some of the approximate methods are quasi-exact in the sense that they can be systematically improved by increasing some accuracy parameter or by accumulating statistics longer; these include the density-matrix renormalization group [41–44], other tensor network states [45], quantum Monte Carlo methods such as FCIQMC [46–48], and coupled cluster (CC) [49]. While some of these approaches have the virtue of even scaling polynomially in the desired accuracy, none of these approaches scale polynomially in molecule size for the generic molecular system. Other, more widely used approximate methods cannot be systematically refined, such as density functional theory (DFT) [50, 51], but permit the study of much larger molecules of up to thousands of atoms.

III. QUANTUM FULL CONFIGURATION INTERACTION ALGORITHM

We now outline the algorithm we use to determine the ground state energy of a small molecule. This algorithm has been previously described in several papers [25]. The first step is to prepare the qubits into a state $|\psi\rangle$ which is a good approximation to the ground state $|\psi_0\rangle$, e.g. it has sufficiently high overlap $\langle\psi|\psi_0\rangle$. For a sufficiently small molecule – like the water molecule used here – this can just be the Hartree-Fock solution $|\psi_{\text{HF}}\rangle$. By choosing as basis functions the single-particle wave functions obtained in a Hartree-Fock calculation we can write $|\psi_{\text{HF}}\rangle = \prod_{i=1}^{N_e} c_i^\dagger |0\rangle$, where $|0\rangle$ is the vacuum, N_e the total number of electrons, and c_i^\dagger creates an electron in the i ’th single-particle state.

For larger molecules we expect the overlap $\langle\psi_{\text{HF}}|\psi_0\rangle$ to decrease significantly. In that case a better approximation to $|\psi_0\rangle$ can be obtained by adiabatic evolution of the wave function from $|\psi_{\text{HF}}\rangle$ towards the true ground

state $|\psi_0\rangle$. This can be achieved by evolving the wave function under the action of a Hamiltonian which slowly evolves from the Hartree-Fock Hamiltonian H_{HF} to the full Coulomb Hamiltonian (1) [22]. Here, the initial Hamiltonian must meet the requirement that it has the Hartree-Fock state as its unique ground state.

In the second part of the QFCI algorithm, the energy of the state $|\psi\rangle$ obtained through the above preparation procedure is measured using the *quantum phase estimation* (QPE) algorithm [2, 52]. This also collapses the state $|\psi\rangle$ (with probability proportional to $|\langle\psi|\psi_0\rangle|$) to the ground state wave-function. We note that there are proposals to improve the quantum phase estimation algorithm [53] as well as proposals to avoid it altogether in the quantum chemistry context [32].

At the highest level, quantum phase estimation takes a state $|\psi\rangle = \sum_i c_i |\phi_i\rangle|0\rangle$, where $|0\rangle$ is the initial state of a number of auxiliary qubits, and converts it into the state $\sum_i c_i |\phi_i\rangle|E_i\rangle$, where $|E_i\rangle$ denotes that a binary representation of the energy has been encoded into the auxiliary qubits. It is very important to note that at the core of the algorithm lies the time evolution of a quantum state, i.e. performing $|\psi(t)\rangle = \exp(-iHt)|\psi(0)\rangle$. The time T required to resolve an absolute error in the energy ϵ is π/ϵ .

On a general-purpose quantum computer, time evolution must be implemented through a circuit composed of one and two-qubit gates. In certain algorithms, most notably Shor’s algorithm, the time evolution $\exp(-iHt)$ can be implemented efficiently by exploiting special properties of the evolution operator, such that the computation time necessary for the whole time evolution scales as $\mathcal{O}(\log t)$. In general, however, a different approach must be taken, the most common one being a Trotter decomposition; other approaches exist, however [30, 32, 36]. In a Trotter decomposition [54, 55], the full time t is divided into discrete time intervals $\Delta_t = t/M$. This leads to a scaling that is at least linear in t , and it incurs a discretization error that is polynomial in Δ_t , with the exponent depending on the type of Trotter decomposition used.

For a Hamiltonian H which is given as a sum over individual terms h_k , the first order Trotter decomposition with M Trotter steps reads

$$\exp(-it \sum_k h_k) = U(\Delta_t)^M + \mathcal{O}(\Delta_t) \quad (2)$$

$$\text{with } U(\Delta_t) = \prod_k \exp(-i\Delta_t h_k). \quad (3)$$

A representation of the Hamiltonian must be chosen where each term $U_k(\Delta_t) = \exp(-i\Delta_t h_k)$ can be broken down into a sequence of standard gates; this is the case for example if each h_k is a product of Pauli matrices. At the cost of a factor of two in the number of circuit elements, a second order Trotter decomposition can be used, which improves the error to $\mathcal{O}(\Delta_t^2)$. In principle, even higher order Trotter breakups can further attenuate

the error. Notice that the time step error coming from using phase estimation with an approximate time evolution operator U is exactly equivalent to an error-free approach with the effective Hamiltonian $H_{\text{eff}} = \ln U(\Delta_t)/(-i\Delta_t)$. The circuit diagrams for all terms that occur in the quantum chemistry Hamiltonian (1) have been previously obtained [25], but to make this presentation self-contained are shown in Appendix D.

In the rest of this paper, we focus solely on the second part of the QFCI algorithm, i.e. measurement of the energy through quantum phase estimation. To compute the computational effort involved in this algorithm, the three factors that need to be taken into account are the number of gates per Trotter step N_g , the total time T that the time evolution must evolve in phase estimation, and the total number of Trotter steps $1/\Delta_t$ needed for evolving for fixed time at fixed error. The total complexity is then $N_g T/\Delta_t$. The total time T is set by the absolute accuracy required. Using an absolute accuracy of 1 milli-Hartree, we get that $T \approx 6000 E_h^{-1}$ is required.

We proceed by computing the Δ_t and N_g required for water and then extrapolate to classically-intractable systems.

IV. RESULTS – WATER

We implement the quantum phase estimation algorithm for a water molecule in a minimal STO-3G basis of ten electrons in fourteen spin-orbitals. We first perform a Hartree-Fock calculation using the PyQuante package [56] and use the thus obtained orthonormalized single particle wave functions as the basis used in the quantum algorithm. We validate our implementation of the quantum algorithm by comparing to full-configuration-interaction (FCI) calculations for the same problem.

Our simulations are performed using the LIQUi|) quantum simulation platform [57], which is an advanced software package developed by Microsoft Research to allow efficient simulation of large quantum circuits (more than 1 million gates) with moderate numbers of qubits (typically 30 qubits in 32GB of memory for Hamiltonian simulations). The system is implemented as an extension to the F# functional programming language [58] and compiles high-level circuit descriptions into targeted simulators (Universal, Stabilizer and Hamiltonian) in one of several environments (Client, Service or Cloud). The architecture is modular and includes packages for optimization, noise modeling, physical gate replacements [59], export and automatic circuit drawing [60]. For this work, a module was added to convert terms into individual circuits, analyze and optimize rotations and then compile re-written circuits into unitary matrices for analysis and simulation.

In Figure 1, we show the energy of a water molecule as a function of bond length and bond angle as obtained from our simulated QFCI calculation. Figure 2 shows the dependence of the accuracy compared to an exact




FIG. 1. (Color online) This figure shows the energy of the water molecule as a function of bond angle and bond length for an STO-3G basis obtained from a restricted HF calculation.




FIG. 2. Discretization error due to the Trotter decomposition for a water molecule at bond length $a = 0.957213$ Å and bond angle $\theta = 104.5225^\circ$.

solution of the same problem on the Trotter time step. We find that a Trotter time step of $\Delta_t = 0.01 E_h^{-1}$ is required to achieve chemical accuracy. We find that, for most orderings of the terms in the Hamiltonian, the error behaves as $\mathcal{O}(\Delta_t^2)$ even for the first-order Trotter decomposition of Eqn. (2). We attribute this to a cancellation in the errors, which seems to be fairly generic. Also, in this regime it does not seem significantly advantageous to go to a higher-order Trotter decomposition. Considering just the scaling of the error, choosing a fourth-order decomposition would allow us to take $\Delta_t = 0.1 E_h^{-1}$ and would thus lead to a ten-fold decrease in the number of Trotter steps, but at the same time would increase the number of gates for a single Trotter step by a larger factor [61]. This trade-off may change in different parameter regimes, for example if a much smaller Δ_t is required. The convergence of our results to the FCI values confirms the correctness of our implementation of simulated QFCI and goes a step beyond the simulations performed previously [22, 25, 27, 31].

We now turn to estimating the resources for this simulation on a quantum computer. To this end, we quan-

tify both the total number of circuit elements and what we refer to as the *parallel depth* of a circuit, which is the minimum required depth of the circuit when exploring possible parallelism between parts of the circuit that operate on disjoint sets of qubits and can hence be executed simultaneously. We also separately count the rotations required as these are the most costly operations for many physical realizations of a quantum computer. We can group the terms in the Hamiltonian into different categories, each requiring a different number of elementary gates. For a fermionic problem, Jordan-Wigner strings [62] are generally used to enforce fermionic signs. These will increase the number of gates necessary to apply an off-diagonal term by a factor of N . Recently, methods have been developed that can improve this [31, 37, 63]; for example, Ref. [37] describes a method where Jordan-Wigner strings can be applied in constant time at the cost of N additional teleportations, which however can be carried out in parallel. It is at this point unclear whether the teleportation can be executed at a similar clock rate as gate operations; nevertheless, we include this possible improvement in our gate counts for parallel operations.

Using the actual gate counts for each term, shown in Table III in Appendix D, we calculate the circuit depth for one Trotter time step and show it in Table II. We find that for our water simulation, the gate count is 20494 for sequential operations and 6438 for parallel operations. Given the required time step $\Delta t = 0.01 E_h^{-1}$ and total time $T = 6 \cdot 10^3 E_h^{-1}$ we must evolve for the QPE to achieve chemical accuracy, which we have established through our simulations above, we need approximately $6 \cdot 10^5$ Trotter steps and 10^{10} serial gates. In this case, working in parallel leads to a speedup of a factor of 3.

V. SCALING TO LARGER MOLECULES

Having set a baseline using our simulations of the water molecule, we need to examine the scaling of two key quantities to be able to extrapolate our results to larger molecules: (i) the number of gates N_g needed for a single Trotter step, (ii) the value of the Trotter time step Δt .

We anticipate the number of gates to be proportional to the number of terms in the Hamiltonian, $\mathcal{O}(N^4)$, multiplied by the number of gates for each term, which due to the Jordan-Wigner transformation is roughly N , thereby giving a scaling of $\mathcal{O}(N^5)$. Note that this assumes that all gates have roughly the same cost; in reality, this may be drastically different. In particular, the number of rotations, which are likely the most expensive gate, does not depend on the Jordan-Wigner strings and will therefore be at most $\mathcal{O}(N^4)$. Also, as mentioned before, there have been proposal to reduce the cost of Jordan-Wigner strings from $\mathcal{O}(N)$ to $\mathcal{O}(1)$ [31, 37, 63]. In Fig. 3 and Table II, we show the scaling of the number of rotations as well as the number of gates in parallel and sequential operation. We find that the empirical scaling matches




FIG. 3. (Color online) (i) Number of rotations, and number of gates in (ii) parallel and (iii) sequential mode. Dashed and dotted lines indicate the expected scaling, which is $(N_{SO})^4$ (dotted lines) for the number of rotations and the parallel gate count, and $(N_{SO})^5$ (dashed line) for the sequential gate count. This data is also listed in Table II.




FIG. 4. (Color online) Scaling of the number of necessary Trotter steps, $1/\Delta t$, with the number of terms in the Hamiltonian m . This is extracted from a series of artificial, but statistically appropriate Hamiltonians with 12, 16, 18, 20 and 24 spin-orbitals. The two curves correspond to different values of the inverse filling $r = N_{SO}/N_e$. The dashed lines indicate fits to $N_{\text{step}} \sim m^\alpha$, with exponents $\alpha(r=3) = 1.27$ and $\alpha(r=2) = 1.08$.

our expectations quite accurately.

Secondly, we need to establish how the Trotter step necessary to attain a given accuracy scales with the system size. Theoretical calculations bound the number of Trotter steps required for a fixed time T by $m^{1+1/2k}$ [61], where m is the number of separate terms the Hamiltonian $H = \sum_{i=1}^m H_i$ is split into, and $2k$ is the order of the Trotter decomposition; see Appendix B for the error in ground state energy due to Trotter error. This theorem

	N_t	Gates/term	$1/\Delta_t$	Total (Parallel)
Upper bound	N^4	N	$(N_t)^{3/2} = N^6$	N^{11} (N^{10})
Empirical scaling	$N^{3.8}$	N	N^4	N^9 (N^8)

TABLE I. Summary of the different contributions to the overall scaling. Here, N_t is the number of terms; gates/term denotes the number of gates required to execute a term using a basic Jordan-Wigner transformation, and assuming all gates take an equal amount of time; Δ_t is the Trotter time step.

of Ref. 61 is stated in terms of the operator norm of H (written as $\|H\|$) which would depend upon N ; however it seems from the derivation that the bound in fact depends upon $\|H_i\|$ which for these synthetic molecules is chosen independent of N . In our case, we effectively use a second-order decomposition and have $\mathcal{O}(N^4)$ terms, such that the number of Trotter steps required is bounded by $\mathcal{O}(N^6)$. Since this is only a bound and since we are not necessarily in the asymptotic regime, it is important to test the scaling empirically. We do this by producing a series of artificial molecules whose Hamiltonian terms have the same statistical properties as real molecules and measuring the scaling as a function of the number of terms in these molecules. Our result is shown in Fig. 4; details are discussed in Appendix A. By performing a fit to the data (shown as dashed line in the figure), we obtained exponents of 1.27 or 1.08, depending on the ratio of the number of electrons to the number of spin orbitals. Hence, we find that the scaling is closer to $N_{\text{step}} = 1/\Delta_t \sim m$ instead of $m^{3/2}$, and we have $1/\Delta_t \sim N^4$.

One possible reason for this improvement is that each term in the Hamiltonian has a nonvanishing commutator with at most $\mathcal{O}(N^3)$ terms, which allows us to give an improved bound on the number of Trotter steps required. See Appendix B for this bound.

We have summarized all these contributions in Table I. We find that the total scaling goes approximately as $\mathcal{O}(N^9)$ for sequential operation. Note that this is better than expected from the most general bounds, which would give a scaling of $\mathcal{O}(N^{11})$. As an example, we take Fe_2S_2 in the STO-3G basis; this molecule has been considered a benchmark example and has enormous importance to biochemistry. Given that the basis of 112 spin orbitals is about 8 times larger than the basis for H_2O , we see that we pay a factor of 8^5 from the number of gates per Trotter step, and a factor of 8^4 from the smaller Trotter step required, leading to a total increase of the runtime by a factor of approximately 10^8 ; we thus have to execute a total number of 10^{18} gates. For Fe_2S_2 , working in parallel (including the execution of Jordan-Wigner strings in constant time) gains us a factor of 20 leaving us with a parallel gate depth of 10^{17} . An additional factor may be gained by choosing a higher-order Trotter decomposition, which may be optimal for larger molecules [61].

A. Quantum hardware requirements for classically-intractable molecules

From the perspective of the number of qubits, the hardware required to simulate a classically-intractable molecule is within reach: current technology can operate on about a dozen qubits, which is about an order of magnitude less than the 100 qubits that are required for the Fe_2S_2 molecule that we have used as an example above. However, achieving the necessary gate count seems much more challenging: current technology allows control for a few hundred gate operations, while we have obtained a gate count of 10^{18} as an upper bound to the number of gates required to simulate Fe_2S_2 . An improvement of more than 15 orders of magnitude over this bound is therefore necessary! Indeed, even the calculation for the ground state of water, which takes mere seconds on a classical computer and requires only few more than a dozen qubits, is many orders of magnitude too large in terms of the gate-count. Even if we assume that Moore’s law applies to quantum computers, i.e. the number of gates that can be executed doubles every 18 months, it will take 75 years to be able to simulate Fe_2S_2 using the QFCI algorithm as discussed here. Even then, assuming a clock speed of 1 GHz (i.e. 1 ns per gate) for gate operations, the calculation for Fe_2S_2 will require 1.5 years to complete! We note that these estimates require the quantum state to be coherent for the entire calculation as the no-cloning theorem forbids checkpointing intermediate results.

Moreover, the numbers we have cited so far are given in terms of *logical* qubits and gate operation, i.e. ideal qubits and gates that are perfectly coherent and operate with perfect fidelity. In reality, qubits have a finite coherence time and gates can only be executed with less than perfect fidelity. Therefore, in most physical implementations, a logical qubit will have to be represented through a number of *physical*, i.e. hardware qubits, and a logical gate operation is obtained via a series of physical gate operations. This allows the use of error correction. In conventional error correction schemes, depending on the fidelity with which the physical gate operations can be executed, between 100 and 10^5 physical gate operations are required for a single logical gate operation; however, it is possible to trade off the number of gates with the number of extra qubits required [64–66]. The alternative route of topological quantum computing [67], where the physical realization of the qubit itself is robust against errors, would require braiding, i.e. adiabatically moving excitations of the underlying topological phase, to operate at least an order of magnitude faster than the required logical gate clock rate. We have also ignored the computation time involved in preparing a state with overlap better than the Hartree-Fock state and the possibility that the whole algorithm might need to be repeated many times to accumulate statistics.

In current technology, typical gate times are $10\mu\text{s}$ for ion traps and 100 ns (10 megahertz) for superconducting

qubits. We note also that the speed at which a quantum computer can run is bounded by the speed at which we can do classical control and it is hard to see going beyond tens of gigahertz in the foreseeable future.

We thus conclude that the simulation of a molecule like Fe_2S_2 on a quantum computer using the QFCI algorithm described in this paper will not be possible by hardware improvements alone; instead, algorithmic improvements of several orders of magnitude will be necessary.

VI. CONCLUSIONS

In this work, we have answered the question of whether a small quantum computer with on the order of 100 qubits will be able to address challenging problems in quantum chemistry that are beyond the reach of classical algorithms. From a purely conceptual point of view, it seems very likely that quantum computers can achieve this: such a quantum computer could trivially represent the wave function of a molecule with up to 100 spin orbitals, and an algorithm (referred to here as QFCI algorithm) is known that in principle should allow the calculation of the ground state energy in polynomial time. Exploring the details of this algorithm, however, we have been able to demonstrate that its polynomial scaling is very large, and that the prefactors work out such that, under reasonable assumptions about improvements in quantum computing, the classically-intractable regime remains intractable also for a quantum computer.

We feel that these estimates draw a line in the sand setting an important barrier that must be overcome for the dream of useful quantum computation for the electronic structure problem in molecules to become a reality. Given even very optimistic assumptions, it seems unlikely that the challenges presented in this paper can be overcome by hardware improvements. Instead, we feel that this emphasizes the importance of *quantum software engineering*, i.e. the necessity of algorithmic advances with a strong focus on practically applicable algorithms. In the realm of classical algorithms, challenging problems are becoming tractable not only due to the fast increase of computational power, but even more so due to advances in algorithms. A prominent example of such an advance are Monte Carlo algorithms with non-local updates, which have lead to performance advances over the original Metropolis algorithm of many orders of magnitude. We believe that similar advances in quantum algorithms will ultimately bring to fruition Feynman’s intuition [11] that quantum computers should be better than classical computers at simulating the properties of quantum systems.

A. The Path Forward

Having focused on the standard QFCI algorithm for solving the QC problem in a given basis of molecular

orbitals, we can ask whether there exists alternate algorithms which might perform better in this regime. For some alternative approaches, see also Refs. 32 and 36. There has been significant work in the literature discussing the more general sparse Hamiltonian problem where algorithms are designed to time-evolve (otherwise structure-less) sparse Hamiltonians. The two current algorithms which scale best (in an incomparable way) are those of Refs. 68 and 69. In the former work, an algorithm is given that scales with the number of non-zero elements per column d and the total time t as $\mathcal{O}(d^2 t \log^3(dt))$, with a complexity polynomial in the logarithm of the inverse error. In our quantum chemistry Hamiltonian, $d = N^2 n_e^2$ where n_e is the number of electrons. Assuming n_e scale as N , we get $\mathcal{O}(N^8 \log^3 N)$. This is asymptotically better than the naive Trotter decomposition used in our QFCI algorithm, with a crossover due to the $\log^3 N$ contribution which is naively at $N \approx 100$. However, the constants may be very different and the bound may not be tight, such that this algorithm may or may not be better than our algorithm discussed in this paper. In the latter work of Ref. [69], the authors use a quantum walk approach to simulate the Hamiltonian and obtain a scaling of $\mathcal{O}(d^{2/3}((\log \log d)t\|H\|)^{4/3})$ with bounded error, where d is the number of non-vanishing elements per row of the matrix. In our problem, $d = N^4$ and therefore we obtain the scaling $\mathcal{O}(N^{8/3}((\log \log N^4)\|H\|)^{4/3})$. Here, $\|H\|$ is the operator norm of the Hamiltonian, which scales at least linearly with the number of electrons (giving at least an N^4 total scaling), but may scale faster if correlation energies dominate. Alternately, Ref. [69] gives a quantum walk algorithm requiring time $\mathcal{O}(d\Lambda_{max}t)$, where Λ_{max} is the largest matrix element of the Hamiltonian. This gives at least N^4 scaling also if Λ_{max} is independent of N . Both of these algorithms require oracle access to matrix elements, which may incur an additional factor of up to N^4 in gate count to encode the coefficients of the Hamiltonians; some of these gates can be executed in parallel, at the cost of additional qubits. It remains to be explored whether this algorithm can be used favorably in some parameter regime for electronic structure calculations. Ongoing work is exploring how these tradeoffs work out in the relevant regime of classically-intractable molecules.

Beyond these, or other new algorithms, we can examine whether potential incremental improvements might chip away at the factors described in this work gaining enough factors of N to make a variant of QFCI tractable.

One problem with the current approach is the need to take *many* time-steps. To decrease this number, a higher-order Trotter decomposition is often suggested. Naively, this is problematic as it comes with significant overhead, but recent work [70, 71] has explored approaches to mitigate this. From a theoretical perspective, even an ‘arbitrarily high-order’ Trotter decomposition changes the scaling with the number of terms m in the decomposition of the Hamiltonian from $m^{1+1/2k}$ to m . As we

are already empirically seeing (and accounting in our estimate) a scaling of m , it is unclear if a higher order Trotter decompositions will garner significant gains; this said, it is possible that these two effects would combine potentially saving a factor of order N^2 . This assumes improvements make these ‘higher order’ decompositions as cheap as a second order decomposition. In a similar vein, it may be possible to extrapolate the time step error or have it cancel out in observables of interest giving us the ability to work at much larger time steps. Finally, we propose an adaptive trotter scheme in Appendix C which might require significantly less gates per time step in certain regimes.

Another possibility is to decrease the number of terms m in the decomposition of the Hamiltonian. For an example of where this has been done, see Ref. [72]. Alternately, a different basis may be used. In a local basis the total number of terms often scale as N^2 instead of N^4 . This would change the scaling of the complete algorithm from $\mathcal{O}(N^9)$ to $\mathcal{O}(N^5)$. In the extreme limit of a real-space basis, there may be significant additional gains coming from the fact that the Hamiltonian can be decomposed into only two pieces [23]. Working in real space brings its own independent set of problems, though, and we are currently looking into whether this is a superior approach in the classically-intractable regime.

ACKNOWLEDGMENTS

We acknowledge useful discussions with M. Reiher, K. Svore, M. Freedman, A. Childs and the participants of the workshop *Quantum Computing for Quantum Chemistry* at Microsoft Research, Redmond, held in November 2012. We acknowledge hospitality of the Aspen Center for Physics, supported by NSF grant 1066293. This research was supported in part by the National Science Foundation under Grant No. NSF PHY11-25915.

After completion of this work, improvements on the methods and bounds discussed here have been proposed in Refs. 76–78.

Appendix A: Scaling of Trotter time step with molecule size

A key component of the overall scaling of the QFCI algorithm comes from the Trotter time step. To keep a fixed accuracy, the Trotter time step will generally have to decrease as the molecule, and hence the number of terms, grows. In this appendix, we will describe our method used to estimate this scaling.

For a given molecule, we estimate the Trotter error by finding the eigenvalues λ_i of the unitary matrix U generated by the time evolution with a given Trotter decomposition for a Trotter timestep Δ_t . We can obtain an estimate for the energies from $\log(\lambda_i)/(-i\Delta_t)$; these are exactly the energy eigenvalues that will be measured

Molecule	Basis	Spin orbitals	Basis size	Sequential		Parallel	
				Rotations	Total	Rotations	Total
H ₂ O	STO-3G	14	441	1.62×10^3	2.05×10^4	1.62×10^3	6.44×10^3
	3-21G	26	1.66×10^6	1.63×10^4	3.81×10^5	1.63×10^4	7.25×10^4
	DZVP	38	1.35×10^8	6.39×10^4	2.12×10^6	6.39×10^4	2.95×10^5
	6-31G**	50	2.8×10^9	1.77×10^5	7.69×10^6	1.77×10^5	8.37×10^5
	6-31++G**	54	6.5×10^9	2.44×10^5	1.14×10^7	2.44×10^5	1.16×10^6
	6-311G**	62	2.9×10^{10}	4.19×10^5	2.23×10^7	4.19×10^5	2.00×10^6
CO ₂	STO-3G	30	1.86×10^6	1.96×10^4	5.32×10^5	1.96×10^4	8.91×10^4
	3-21G	54	1.70×10^{14}	1.34×10^5	6.19×10^6	1.34×10^5	6.30×10^5
	DZVP	90	1.03×10^{20}	1.39×10^6	1.06×10^8	1.39×10^6	6.79×10^6
	6-311G**	90	1.03×10^{20}	1.02×10^6	7.62×10^7	1.02×10^6	4.95×10^6
Fe ₂ S ₂	STO-3G	112	3.4×10^{25}	7.44×10^6	6.31×10^8	7.44×10^6	3.59×10^7
	3-21G	168	2.8×10^{48}	4.51×10^7	5.84×10^{10}	4.51×10^7	2.21×10^8

TABLE II. Gate count for one Trotter step using a sequential or a parallel circuit. The basis sizes are extracted from the PyQuante package [56]; for details on the basis sets, see also Refs. [73–75].

with quantum phase estimation. To perform this diagonalization for large enough molecules, we need to resort to iterative diagonalization techniques such as the Arnoldi method. These methods generally only extract a few eigenvalues; to target the eigenvalues of U corresponding to the ground state, we shift the Hamiltonian by an appropriate amount ϵ such that $E_0 + \epsilon > 0$; for small enough Δ_t , the eigenvalue we are interested in is then the eigenvalue closest to 1 on the unit circle, and hence the eigenvalue with the largest real component.

In order to extract the scaling, we need to perform estimates for a number of different molecules; however, the exponential scaling of the classical algorithm with the number of spin orbitals as well as the large overhead of estimating the energy levels from U instead of a direct calculation of the ground state, as would usually be done in FCI, severely restrict the number of spin orbitals we can study to roughly $N = 24$. In order to have a sufficient number of generic molecules in the range $N = 8$ to $N = 24$, we generate random Hamiltonians that imitate the statistics of interaction terms found in real molecules. Specifically, we generate terms $h_{pp}c_p^\dagger c_p$, $h_{pq}c_p^\dagger c_q$, $h_{pqqp}c_p^\dagger c_q^\dagger c_q c_p$, $h_{pqqr}c_p^\dagger c_q^\dagger c_q c_r$ and $h_{pqrs}c_p^\dagger c_q^\dagger c_r c_s$ with the distribution functions for the parameters chosen as (here $u_{[a,b]}(x)$ is a uniform distribution of values between a and b evaluated at x):

$$p(h_{pp}) = u_{[-10,0]}(h_{pp}) \quad (\text{A1})$$

$$p(h_{pq}) = u_{[-1,1]}(h_{pq}) \quad (\text{A2})$$

$$p(h_{pqqp}) = u_{[-0.5,0.5]}(h_{pqqp}) \quad (\text{A3})$$

$$p(h_{pqqr}) = \frac{1}{2 \cdot 0.2} e^{-|h_{pqqr}|/0.2} \quad (\text{A4})$$

$$p(h_{pqrs}) = \frac{1}{2 \cdot 0.1} e^{-|h_{pqrs}|/0.1} \quad (\text{A5})$$

We only keep terms compatible with particle-number conservation symmetry; additionally, we remove a fraction of the terms to mirror the fact that in a real molecule, terms may be forbidden by spatial symmetries. We generate molecules up to $N = 24$, obtain their ground state

energy for $N_e = N/2$ and $N_e = N/3$, as well as the respective energy estimates for a number of different values of Δ_t , to obtain the error estimate $\epsilon(\Delta_t)$. The number of Trotter time steps required to reach accuracy ϵ_t is then given as

$$N_{\text{Trotter}} = \left(\frac{\epsilon(\Delta_t)}{\Delta_t^2 \epsilon_t} \right)^{1/2} \quad (\text{A6})$$

assuming that $\epsilon(\Delta_t) \sim \Delta_t^2$, which we empirically confirm.

Appendix B: Improved Trotter-Suzuki Error Bounds and Ground State Energy Error

Given a Hamiltonian $H = \sum_{i=1}^m H_i$, the second order Trotter-Suzuki approximation to $\exp(-iH\Delta_t)$ is given by

$$U^{TS} \equiv \exp(-iH_1\Delta_t/2) \dots \exp(-iH_m\Delta_t/2) \quad (\text{B1}) \\ \times \exp(-iH_m\Delta_t/2) \dots \exp(-iH_1\Delta_t/2).$$

The usual derivation of the bound on second order Trotter-Suzuki error proceeds by expanding both the exact expression $U = \exp(-iH\Delta_t)$ and the second order Trotter-Suzuki approximation U^{TS} to third order in Δ_t and showing that they agree at first and second order and bounding the error at third order. The error is shown to be bounded by $(m\Lambda\Delta_t)^3$, where Λ is an upper bound to $\|H_i\|$. Thus, for evolution for a fixed time t , with time step Δ_t , the error that accumulates is at most $t(m\Lambda)^3\Delta_t^2$, meaning that to have small error at fixed t, Λ it suffices to have $\Delta_t \ll m^{-3/2}$. The Trotter number scales as $1/\Delta_t$ so one needs a Trotter number scaling as $m^{3/2}$.

Before giving an improved bound, let us consider the impact of this Trotter-Suzuki error on the ground state energy. The bounds on Trotter-Suzuki error are typically expressed by stating an error estimate of the form $\|U - U^{TS}\| \leq \epsilon$, for some number ϵ . If this is the error for a single time-step Δ_t , then the error for evolution over a large time t may be much larger: $\|U^{t/\Delta_t} - (U^{TS})^{t/\Delta_t}\| \leq$

$(t/\Delta_t)\epsilon$. Hence, it might seem that since phase estimation requires evolution for $t \gg \Delta_t$, we will encounter a large error. In fact, this is not true. Phase estimation does require evolution for large time $t \gg \Delta_t$, but this is done simply to estimate eigenvalues of U^{TS} . Hence, the error in energy can be directly obtained from the difference between eigenvalues of U and U^{TS} . We will assume that the time step is short enough that all eigenvalues of both U and U^{TS} lie sufficiently close to 1 that there is no ambiguity in which branch cut of the logarithm should be used to determine the energy. So, the error in ground state energy is bounded by

$$\frac{1}{\Delta_t} \|U - U^{TS}\|,$$

and so for second order Trotter-Suzuki we will obtain an error in ground state energy that scales as $(m\Lambda^3)\Delta_t^2$ at worst.

However, in a quantum chemistry setting, we have the case that many of the commutators $[H_i, H_j]$ are equal to zero. This significantly reduces the error. Let K be the maximum over i of the number of terms H_j which have a nonvanishing commutator with H_i . For the Hamiltonian (1), $K = \mathcal{O}(N^3)$, since any two terms that do not commute must agree on at least one index. As a result, we can bound the error by

$$\|U - U^{TS}\| \leq \mathcal{O}(mK^2\Lambda^3\Delta_t^3), \quad (\text{B2})$$

rather than the previous bound of $\mathcal{O}(m^3\Lambda^3\Delta_t^3)$, as we now show. For our problem, with $m = \mathcal{O}(N^4)$, $K = \mathcal{O}(N^3)$, this means that a Trotter number of $\mathcal{O}(N^5)$ suffices to obtain small error.

We begin with a slightly naive derivation of the bound above using series expansions. This bound suffers from some problems as we point out at the end, and afterwards we give a corrected proof of the result. The series expansion up to third order can be written as a sum of many terms. All terms cancel at first and second order. There are terms which are third order in H_i for some given i ; these also cancel exactly. There are terms which are first order in H_i and second order in H_j for some $j \neq i$. There are at most mk such terms, so they contribute at most $\mathcal{O}(mk\Delta^3\Delta_t^3)$ to the total error. Finally, there are terms which are first order in H_i, H_j, H_k for $i \neq j \neq k \neq i$. However these terms also cancel if any one of the three operators H_i, H_j, H_k commutes with the other two operators. Thus, there are only at most $\mathcal{O}(mK^2\Delta_t^3)$ such terms. Thus, the third order error is bounded by $\mathcal{O}(mK^2\Delta_t^3)$.

A similar calculation can be done at higher order. In general, it is useful to introduce a ‘‘linked cluster expansion’’ to keep track of combinatorics. Consider a term at q -th order which is linear in $H_{i_1}, H_{i_2}, \dots, H_{i_m}$, where the i_a need not be distinct from each other. We introduce a diagrammatic notation, writing q distinct points corresponding to the terms and drawing a line between the a -th points and the b -th point if $[H_{i_a}, H_{i_b}] \neq 0$. De-

fine the ‘‘linked clusters’’ to be the connected components of the resulting graph. A term at m -th order cancels unless at least one of the clusters contains at least three points. Summing over all clusters satisfying this condition, this bounds the expression at q -th order by $\mathcal{O}(m^{q-2}K^2(\Lambda\Delta_t)^q)$.

The trouble with this series expansion method is twofold. First, there is a combinatoric issue of bounding the prefactors in front of the higher order terms. While this can be dealt with, the more serious issue is that the expansion parameter in the series expansion is actually still $m\Lambda\Delta_t$, so the bound $\mathcal{O}(mK^2\Lambda^3\Delta_t^3)$ can only be proven for sufficiently small $m\Lambda\Delta_t$.

We now give an alternate derivation that does not assume $m\Lambda\Delta_t$ is small also. We first bound the error in the second order Suzuki expansion for a problem with only two terms, called A and B . Let us fix $\Delta_t = 1$. Define $H(x) = B + (1-x)A$. We wish to bound

$$\begin{aligned} & \left\| \exp(-i\frac{A}{2}) \exp(-iB) \exp(-i\frac{A}{2}) - \exp(-iH(0)) \right\| \\ &= \left\| \int_0^1 \partial_x \left(\exp(-ix\frac{A}{2}) \exp(-iH(x)) \exp(-ix\frac{A}{2}) \right) dx \right\|. \end{aligned}$$

We will bound the norm of the derivative on the right-hand side of the above expression, and then integrate this bound over $0 \leq x \leq 1$ to bound the second order error. Let $A(t, x) = \exp(iH(x)t)A \exp(-iH(x)t)$. Then,

$$\begin{aligned} & \left\| \partial_x \left(\exp(-ix\frac{A}{2}) \exp(-iH(x)) \exp(-ix\frac{A}{2}) \right) \right\| \\ &= \left\| \exp(-ix\frac{A}{2}) \exp(-iH(x)) \left(-i\frac{A}{2} - i\frac{A(1,x)}{2} \right. \right. \\ & \quad \left. \left. + i \int_0^1 A(t, x) dt \right) \exp(-ix\frac{A}{2}) \right\| \\ &= \left\| -\frac{A(0, x)}{2} - \frac{A(1, x)}{2} + \int_0^1 A(t, x) dt \right\|. \end{aligned}$$

We have $A(t, x) = A(0, x) + tA'(0, x) + \int_0^t (t-s)A''(s, x)ds$, where A' and A'' represent first and second derivatives with respect to t in $A(t, x)$. One may verify that the terms in A and A' cancel in the above equation, leaving only terms in A'' , giving after some calculus

$$\begin{aligned} & \left\| \partial_x \left(\exp(-ix\frac{A}{2}) \exp(-iH(x)) \exp(-ix\frac{A}{2}) \right) \right\| \\ &= \left\| \int_0^1 \frac{s-s^2}{2} A''(s, x) ds \right\| \\ &\leq \int_0^1 \|A''(s, x)\| ds. \end{aligned}$$

Integrating over x , and using $\|A''(s, x)\| = \|[[A, H(x)], H(x)]\| \leq \|[[A, B], A] + [[A, B], B]\|$, we find that

$$\begin{aligned} & \left\| \exp(-i\frac{A}{2}) \exp(-iB) \exp(-i\frac{A}{2}) - \exp(-i(A+B)) \right\| \\ &\leq \|[[A, B], A]\| + \|[[A, B], B]\|. \quad (\text{B3}) \end{aligned}$$

We now apply Eq. (B3) inductively to give the desired bound on $U - U^{TS}$. Let $U_j = \exp(-i \sum_{j \leq k \leq m} H_k)$ and let

$$U_j^{TS} = \exp(-iH_j \frac{\Delta_t}{2}) \dots \exp(-iH_m \frac{\Delta_t}{2}) \\ \times \exp(-iH_m \frac{\Delta_t}{2}) \dots \exp(-iH_j \frac{\Delta_t}{2}). \quad (\text{B4})$$

Then by a triangle inequality,

$$\|U_{j-1} - U_{j-1}^{TS}\| \quad (\text{B5}) \\ \leq \|U_{j-1} - \exp(iH_{j-1} \frac{\Delta_t}{2}) U_j \exp(-iH_{j-1} \frac{\Delta_t}{2})\| \\ + \|\exp(iH_{j-1} \frac{\Delta_t}{2}) U_j \exp(-iH_{j-1} \frac{\Delta_t}{2}) - U_{j-1}^{TS}\| \\ = \|U_{j-1} - \exp(iH_{j-1} \frac{\Delta_t}{2}) U_j \exp(-iH_{j-1} \frac{\Delta_t}{2})\| \\ + \|U_j - U_j^{TS}\|.$$

Using Eq. (B3) to bound the first term on the right-hand side of the above equation, taking $A = H_{j-1} \Delta_t$ and $B = \sum_{j \leq k \leq m} H_k \Delta_t$, we get

$$\|U_{j-1} - U_{j-1}^{TS}\| \leq \text{const.} \times \Lambda^3 \Delta_t^3 K^2 + \|U_j - U_j^{TS}\|, \quad (\text{B6})$$

where the constant is a numeric constant of order unity. Summing over j , we obtain the desired bound Eq. (B2) for $\|U - U^{TS}\| = \|U_1 - U_1^{TS}\|$.

It seems likely that this same proof will also work to bound higher order errors in higher order Trotter-Suzuki expressions. For $2r$ -th order Trotter-Suzuki, we expect to bound the error by $\mathcal{O}(mk^{2r} \Lambda^{2r+1} \Delta_t^{2r+1})$. If indeed this holds, for $K = N^3$ and $m = N^4$, for sixth order Trotter-Suzuki the Trotter number required to obtain given error is actually *sublinear* in m .

Appendix C: Coalescing

Our work makes it clear that algorithmic improvements are needed to make quantum chemistry practical on a quantum computer. Here we suggest one potential improvement which significantly decreases the total number of gate operations required by using an adaptive Trotter decomposition for terms of different magnitude. Unfortunately, there is some tradeoff in this scheme as the rearrangement we suggest likely increases the Trotter error. Further research is required to understand, in what regimes, this tradeoff is such that this is a beneficial approach.

The key idea is based on the observation that in the Hartree-Fock basis, many off-diagonal matrix elements are very small. In the unitary time evolution these small terms can be applied with much larger Trotter time steps Δ_t compared to the larger terms. As a consequence fewer circuits have to be applied and the total gate count is significantly reduced.

Specifically, we propose to use a different time step $\Delta_t^{(k)}$ for each of the terms in the Hamiltonian h_k and

to choose them such that the product of the amplitude of the term and the time step $\Delta_t^{(k)}$. $\|h_k\|$ is roughly homogeneous across all terms.

Equivalently, consider the Trotter-decomposed time evolution operator for a given total time T ,

$$U = \left(\prod U_k \right)^{T/\Delta_t}, \quad (\text{C1})$$

where U_k applies the Hamiltonian term h_k for a Trotter step Δ_t . In the Hartree-Fock basis, many of these terms are extremely close to the identity. We can then imagine rearranging the terms in this expansion such that identical terms that are very close to the identity, i.e. that have a very small coefficient in the Hamiltonian, are grouped together and can be executed at once.

From this description it becomes clear that some additional discretization error will be accumulated by changing the order of terms, and it may therefore be necessary to reduce the Trotter time step to keep the total error fixed. This leads to a trade-off between grouping terms and keeping the error constant. The best scheme within this trade-off depends sensitively on the specific distribution of Hamiltonian terms h_k . In this paper, we do not address this question in detail, but defer it to future work. However, we do give a brief theoretical analysis of one simple coalescing scheme.

Let the Hamiltonian H be a sum of terms $H = \sum_i H_i$. Divide these terms H_i into *buckets*, so that every term is in exactly one bucket. Label the buckets $1, \dots, k$, and let B_a be the set of integers i so that H_i is in bucket a . Assume that there are N_a terms in bucket a with $\|H_i\| \leq \Lambda_a$ for $i \in B_a$. We arrange the buckets to contain terms in decreasing order of magnitude, so that B_1 contains the largest terms and is executed the most frequently, while later buckets are executed less frequently.

We wish to approximate $U = \exp(-iHt)$ with a quantum circuit, and we assume that we have circuits to implement $\exp(-i\theta H_i)$ for any i .

We analyze a coalescing scheme to approximate this unitary, and show that it achieves error ϵ in operator norm bounded by

$$\epsilon \leq \sum_a O\left(\frac{S_a^3 + T_{a-1} S_a^2 + T_{a-1}^2 S_a}{(2^{k-a})^2}\right), \quad (\text{C2})$$

where we define

$$S_a = N_a \Lambda_a t \quad (\text{C3})$$

and

$$T_a = \sum_{b \leq a} S_b, \quad (\text{C4})$$

with $T_0 = 0$. The scheme is defined by Eqs. (C8,C9,C11), with V_k defined by those equations being the approximation to U .

To gain some intuition for Eq. (C2), let $K_a = \sum_{i \in B_a} H_i$ be the sum of all terms in bucket a . Then, the

first term $O(S_a^3/(2^{k-a})^2)$ in the equation is the error we obtain by approximating $\exp(-iK_a t)$ by doing a second order Trotter-Suzuki expansion to $\exp(-iK_a \frac{t}{2^{k-a}})$ and then taking the 2^{k-a} -th power of that approximation. In a sense, this first term results from errors in commutators in a single bucket. The terms in $T_{a-1}S_a^2$ and $T_{a-1}^2S_a$ result from interaction between buckets. The terms resulting from interaction between buckets have one unfortunate effect: if there is a large term in the first bucket (which gives a large T_a and which will be approximated with a very short Trotter step), then ideally we would like to suppress the error by a large denominator $(2^{k-1})^2$, and indeed the first term in Eq. (C2) is suppressed by this factor. However, since this term interacts with terms in later buckets, there will be error terms with denominators of order $(2^{k-a})^2$ for all a also appearing in the second and third terms in Eq. (C2) and this is a smaller denominator. However, these terms will be at most second order in T_1 and will be at least first order in S_a so if S_a is small then this will help suppress these terms.

To define and analyze the scheme, let

$$J_a = \sum_{b \leq a} \sum_{i \in B_b} H_i \quad (\text{C5})$$

be the sum of all terms in buckets $1, \dots, a$ and let

$$U_a = \exp(-iJ_a \frac{t}{2^{k-a}}), \quad (\text{C6})$$

so that $U = U_k$. We proceed inductively, using an approximation to U_a to construct an approximation to U_{a+1} . We call this approximation V_a . and let

$$\epsilon_a = \|V_a - U_a\|. \quad (\text{C7})$$

We define V_1 by a standard second order Trotter-Suzuki as follows. Let

$$W_a = \prod_{i \in B_1} \exp(-i\frac{1}{2}H_i \frac{t}{2^{k-a}}), \quad (\text{C8})$$

where the product is taken in any fixed arbitrary order, and let W'_a denote the same product as W_a except taken in the reverse order (if H is real, then $W'_a = W_a^T$). Then, let

$$V_1 = W_1 W'_1. \quad (\text{C9})$$

By standard estimates,

$$\epsilon_1 \leq O\left(\frac{S_1}{2^{k-1}}\right)^3. \quad (\text{C10})$$

This estimate for ϵ_1 is based on a third order Taylor expansion; the lower order terms in the Taylor expansion cancel, and the higher order terms in the Taylor expansion are higher order in S_1 .

We then define

$$V_{a+1} = W_{a+1} V_a^2 W'_{a+1}. \quad (\text{C11})$$

Note that

$$\|V_{a+1} - W_{a+1} U_a^2 W'_{a+1}\| \leq 2\epsilon_a. \quad (\text{C12})$$

Also,

$$\begin{aligned} & \|U_{a+1} - W_{a+1} U_a^2 W'_{a+1}\| \\ & \leq O\left(\frac{S_{a+1}^3 + \|J_a\| S_{a+1}^2 + \|J_a\|^2 S_{a+1}}{(2^{k-a-1})^3}\right), \end{aligned} \quad (\text{C13})$$

as can be estimated using a Taylor series similar to before. So, by Eqs. (C12, C13),

$$\epsilon_{a+1} \leq 2\epsilon_a + O\left(\frac{S_{a+1}^3 + T_a S_{a+1}^2 + T_a^2 S_{a+1}}{(2^{k-a-1})^3}\right), \quad (\text{C14})$$

using the fact that $\|J_a\| \leq T_a$.

Hence,

$$\epsilon_a \leq \sum_{b \leq a} 2^{a-b} O\left(\frac{S_b^3 + T_{b-1} S_b^2 + T_{b-1}^2 S_b}{(2^{k-b-1})^3}\right), \quad (\text{C15})$$

and

$$\epsilon_k \leq \sum_a O\left(\frac{S_a^3 + T_{a-1} S_a^2 + T_{a-1}^2 S_a}{(2^{k-a})^2}\right). \quad (\text{C16})$$

Appendix D: Circuits

Following the approach set out in Ref. [25], we convert each term in Equation 1 via a Jordan-Wigner transformation which expresses fermionic operators in terms of Pauli spin operators $\{\sigma^x, \sigma^y, \sigma^z, \mathbf{1}\}$. This technique allows us to preserve the necessary commutation relations by entangling all of the spin-orbitals (qubits) between annihilation and creation operators with **CNOT** gates, performing the unitary propagator ($\exp(-iHt)$) at a given strength (t_{pq} or V_{pqrs}) and then unentangling by reversing the order of the applied **CNOTs**. If the propagator is in the computational basis σ^z , then a rotation of the desired strength may be applied directly. If the propagator is in the σ^x basis then the basis must be flipped from σ^x to σ^z with a Hadamard transformation **H**; the operation is then performed and then returned from σ^z to σ^x with an additional **H** at the end. Likewise, **Y** and **Y**[†] may be used to flip between σ^y and σ^z . The basis flip operations are:

$$\mathbf{H} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad \mathbf{Y} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \quad (\text{D1})$$

The resulting circuits are presented in Figs. 5, 6, 7 and 8. These circuits coincide with those of Ref. [25], except for those for H_{pqqr} which were not needed in the example discussed in the reference. In all circuits, θ refers to the term strength (t_{pq} or V_{pqrs}) computed by a standard technique. In our case this was PyQuante [56] running a Restricted Hartree-Fock model.




FIG. 5. Circuit representations of Hamiltonian terms H_{pp} and H_{pq} .




FIG. 6. Circuit representation of Hamiltonian terms H_{pqqp} .

The circuits shown above are easily parameterized to show expected resource usage as much larger molecules

are modeled. Table III shows exact counts for both sequential and parallel implementation of all one and two body terms used in the the Hamiltonian.

-
- [1] P. Shor, in *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on* (1994) pp. 124–134.
- [2] A. Y. Kitaev, [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).
- [3] P. W. Shor, *SIAM journal on computing* **26**, 1484 (1997).
- [4] S. Hallgren, *Journal of the ACM (JACM)* **54**, 4 (2007).
- [5] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, *Phys. Rev. A* **64**, 022319 (2001).
- [6] D. Aharonov, V. Jones, and Z. Landau, *Algorithmica* **55**, 395 (2009).
- [7] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, in *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing* (ACM, 2003) pp. 59–68.
- [8] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, *Nature* **414**, 883 (2001).
- [9] B. P. Lanyon *et al.*, *Nat Chem* **2**, 106 (2010).
- [10] J. Du, N. Xu, X. Peng, P. Wang, S. Wu, and D. Lu, *Phys. Rev. Lett.* **104**, 030502 (2010).
- [11] R. Feynman, *International Journal of Theoretical Physics* **21**, 467 (1982).
- [12] H. Nakano and T. Sakai, *J. Phys. Soc. Jpn.* **80**, 053704 (2011).
- [13] A. M. Läuchli, J. Sudan, and E. S. Sørensen, *Physical Review B* **83**, 212401 (2011).
- [14] S. Capponi, O. Derzhko, A. Honecker, A. M. Läuchli, and J. Richter, *Phys. Rev. B* **88**, 144416 (2013).
- [15] Z. Gan and R. J. Harrison, in *Supercomputing, 2005. Proceedings of the ACM/IEEE SC 2005 Conference* (IEEE, 2005) pp. 22–22.
- [16] Y. Kurashige, G. K.-L. Chan, and T. Yanai, *Nat Chem* **5**, 660 (2013).
- [17] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, *Phys. Rev. Lett.* **106**, 130506 (2011).
- [18] S. Beauregard, *Quantum Info. Comput.* **3**, 175 (2003).
- [19] M. G. Whitney, N. Isailovic, Y. Patel, and J. Kubiatowicz, *ACM SIGARCH Computer Architecture News* **37**, 383 (2009).
- [20] T. Kleinjung, K. Aoki, J. Franke, A. Lenstra, E. Thomé, J. Bos, P. Gaudry, A. Kruppa, P. Montgomery, D. A. Os- vik, H. te Riele, A. Timofeev, and P. Zimmermann, “Factorization of a 768-bit rsa modulus,” Cryptology ePrint Archive, Report 2010/006 (2010), <http://eprint.iacr.org/>.
- [21] D. A. Lidar and H. Wang, *Phys. Rev. E* **59**, 2429 (1999).
- [22] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).
- [23] I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik, *Proceedings of the National Academy of Sciences* **105**, 18681 (2008).
- [24] H. Wang, S. Kais, A. Aspuru-Guzik, and M. R. Hoffmann, *Phys. Chem. Chem. Phys.* **10**, 5388 (2008).






FIG. 7. Circuit representations of Hamiltonian terms (i) H_{pqr} , $p < q < r$, and (ii) H_{pqr} , $q < p$ or $q < r$.

(i)



(ii)

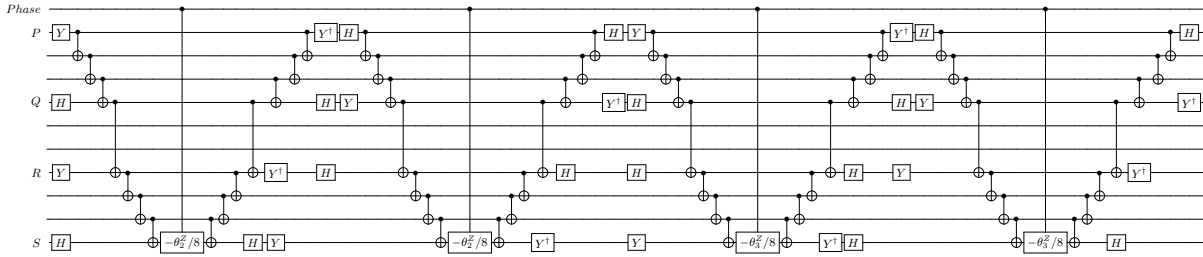


FIG. 8. Circuit representation of Hamiltonian terms H_{pqr} .

Sequential Circuit	Global R_z	H, Y, Y^\dagger	CNOT	CR_z	BSM	Total
H_{pp}	1	8	$4(q-p)$	1		1
H_{pq}				2		$10 + 4(q-p)$
H_{pqqp}				3		$1+5$
$H_{pqqr}, p < q < r$				4		$12 + 4(r-p)$
$H_{pqqr}, q < p$ or $r < q$				4		$16 + 4(r-p)$
H_{pqrs}				4		$8 \cdot 1$
		$8 \cdot 8$	$8 \cdot 2(q-p+s-r+1)$	$8 \cdot 1$		$8 \cdot 9 + 8 \cdot 2(q-p+s-r+1)$
Parallel Circuit	Global R_z	H, Y, Y^\dagger	CNOT	CR_z	BSM	Total
H_{pp}	1	8	4	1	4	1
H_{pq}				2		18
H_{pqqp}				3		$1+5$
H_{pqqr}				4		24
H_{pqrs}				4		$8 \cdot 7$
						$8 \cdot 2$

TABLE III. Circuit depths of the circuits for the individual terms in the Hamiltonian for sequential (top) and parallel (bottom) execution of the circuits. The global rotation gate in the H_{pqqp} term needs to be counted only once for all such terms. The eight terms in H_{pqrs} can be reduced to 4 or 2 for terms where some of the angles θ_i are 0 due to symmetry. For the parallel circuits, the method of Ref. [37] can execute the Jordan-Wigner strings in constant time at the cost of additional Bell-state measurements (BSM).

- [25] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, *Molecular Physics* **109**, 735 (2011).
- [26] I. Kassal and A. Aspuru-Guzik, *The Journal of Chemical Physics* **131**, 224102 (2009).
- [27] L. Veis and J. Pittner, *The Journal of Chemical Physics* **133**, 194106 (2010).
- [28] I. Kassal, J. D. Whitfield, A. Perdomo-Ortiz, M.-H. Yung, and A. Aspuru-Guzik, *Annual Review of Physical Chemistry* **62**, 185 (2011).
- [29] A. Daskin and S. Kais, *The Journal of Chemical Physics* **134**, 144112 (2011).
- [30] A. Daskin, A. Grama, G. Kollias, and S. Kais, *The Journal of Chemical Physics* **137**, 234112 (2012).
- [31] J. T. Seeley, M. J. Richard, and P. J. Love, *The Journal of Chemical Physics* **137**, 224109 (2012).
- [32] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, Preprint (2013), [arXiv:1304.3061](https://arxiv.org/abs/1304.3061).
- [33] J. D. Whitfield, Preprint (2013), [arXiv:1306.1147](https://arxiv.org/abs/1306.1147).
- [34] M.-H. Yung, J. Casanova, A. Mezzacapo, J. McClean, L. Lamata, A. Aspuru-Guzik, and E. Solano, Preprint (2013), [arXiv:1307.4326](https://arxiv.org/abs/1307.4326).
- [35] L. Lamata, A. Mezzacapo, J. Casanova, and E. Solano, Preprint (2013), [arXiv:1312.2849](https://arxiv.org/abs/1312.2849).
- [36] R. Babbush, P. J. Love, and A. Aspuru-Guzik, Preprint (2013), [arXiv:1311.3967](https://arxiv.org/abs/1311.3967).
- [37] N. Cody Jones, J. D. Whitfield, P. L. McMahon, M.-H. Yung, R. V. Meter, A. Aspuru-Guzik, and Y. Yamamoto, *New Journal of Physics* **14**, 115023 (2012).
- [38] S. Kais, ed., *Quantum Information and Computation for Chemistry*, Advances in Chemical Physics, Vol. 154 (Wiley, 2014).
- [39] M.-H. Yung, J. D. Whitfield, S. Boixo, D. G. Tempel, and A. Aspuru-Guzik, Preprint (2012), [arXiv:1203.1331](https://arxiv.org/abs/1203.1331).
- [40] M. Podewitz, M. T. Stiebritz, and M. Reiher, *Faraday Discuss.* **148**, 119 (2011).
- [41] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [42] S. R. White and R. M. Noack, *Phys. Rev. Lett.* **68**, 3487 (1992).
- [43] K. H. Marti and M. Reiher, in *Modern and Universal First-principles Methods for Many-Electron Systems in Chemistry and Physics*, Progress in Physical Chemistry, Vol. 3 (R Oldenbourg Verlag, 2010).
- [44] G. K.-L. Chan and S. Sharma, *Annual Review of Physical Chemistry* **62**, 465 (2011).
- [45] K. H. Marti, B. Bauer, M. Reiher, M. Troyer, and F. Verstraete, *New Journal of Physics* **12**, 103008 (2010).
- [46] G. H. Booth, A. J. W. Thom, and A. Alavi, *The Journal of Chemical Physics* **131**, 054106 (2009).
- [47] D. Cleland, G. H. Booth, and A. Alavi, *The Journal of Chemical Physics* **132**, 041103 (2010).
- [48] G. H. Booth, A. Gruneis, G. Kresse, and A. Alavi, *Nature* **493**, 365 (2013).
- [49] J. Čížek, *The Journal of Chemical Physics* **45**, 4256 (1966).
- [50] P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964).
- [51] W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
- [52] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and quantum computation*, 47 (American Mathematical Soc., 2002).
- [53] K. M. Svore, M. B. Hastings, and M. Freedman, *Quant. Inf. Comp.* **14**, 306 (2013).
- [54] H. F. Trotter, *Proc. Amer. Math. Soc.* **10**, 545 (1959).
- [55] M. Suzuki, *Communications in Mathematical Physics* **51**, 183 (1976).
- [56] Available at <http://pyquante.sourceforge.net>.
- [57] D. Wecker and K. Svore, Submitted (2013).
- [58] D. Syme, A. Granicz, and A. Cisternino, *Expert F#* (Apress, 2007).
- [59] C. M. Dawson and M. A. Nielsen, *Quantum Information & Computation* **6**, 81 (2006), [arXiv:quant-ph/0505030](https://arxiv.org/abs/quant-ph/0505030).
- [60] All circuits in this paper were auto-generated by LIQ*U*i).
- [61] D. Berry, G. Ahokas, R. Cleve, and B. Sanders, *Commun. Math. Phys.* **270**, 359 (2007).
- [62] P. Jordan and E. Wigner, *Z. Phys. A* **47**, 631 (1928).
- [63] S. B. Bravyi and A. Y. Kitaev, *Annals of Physics* **298**, 210 (2002), [quant-ph/0003137](https://arxiv.org/abs/quant-ph/0003137).
- [64] D. Gottesman, in *Quantum Information Science and Its Contributions to Mathematics, Proceedings of Symposia in Applied Mathematics*, Vol. 68 (2009) p. 13.
- [65] S. J. Devitt, W. J. Munro, and K. Nemoto, *Reports on Progress in Physics* **76**, 076001 (2013).
- [66] D. Gottesman, Preprint (2013), [arXiv:1310.2984](https://arxiv.org/abs/1310.2984).

- [67] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, *Rev. Mod. Phys.* **80**, 1083 (2008).
- [68] D. W. Berry, R. Cleve, and R. D. Somma, Preprint (2013), [arXiv:1308.5424](#).
- [69] D. W. Berry and A. M. Childs, *Quantum Information & Computation* **12**, 29 (2012), [arXiv:0910.4157](#).
- [70] A. M. Childs and N. Wiebe, *J. Math. Phys.* **54**, 062202 (2013), [arXiv:1211.4945](#).
- [71] A. M. Childs and N. Wiebe, *Quantum Information & Computation* **12**, 901 (2012), [arXiv:1202.5822](#).
- [72] A. M. Childs and R. Kothari, in *Theory of Quantum Computation, Communication, and Cryptography*, Lecture Notes in Computer Science, Vol. 6519, edited by W. Dam, V. Kendon, and S. Severini (Springer Berlin Heidelberg, 2011) pp. 94–103.
- [73] “Basis set exchange,” Available at <http://bse.pnl.gov/bse/portal>.
- [74] D. Feller, *Journal of computational chemistry* **17**, 1571 (1996).
- [75] K. L. Schuchardt, B. T. Didier, T. Elsethagen, L. Sun, V. Gurumoorthi, J. Chase, J. Li, and T. L. Windus, *Journal of chemical information and modeling* **47**, 1045 (2007).
- [76] B. Toloui and P. J. Love, Preprint (2013), [arXiv:1312.2579](#).
- [77] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, Preprint (2014), [arXiv:1403.1539](#).
- [78] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, Preprint (2014), [arXiv:1406.4920](#).