

A Comparison of the Bravyi–Kitaev and Jordan–Wigner Transformations for the Quantum Simulation of Quantum Chemistry

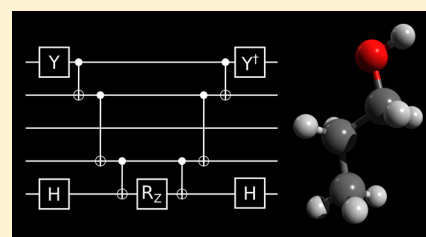
Andrew Tranter,[†] Peter J. Love,[¶] Florian Mintert,[†] and Peter V. Coveney^{*,‡,¶}

[†]Department of Physics, Imperial College London, London SW7 2AZ, United Kingdom

[‡]Centre for Computational Science, University College London, London WC1H 0AJ, United Kingdom

[¶]Department of Physics, Tufts University, Medford, Massachusetts 02155, United States

ABSTRACT: The ability to perform classically intractable electronic structure calculations is often cited as one of the principal applications of quantum computing. A great deal of theoretical algorithmic development has been performed in support of this goal. Most techniques require a scheme for mapping electronic states and operations to states of and operations upon qubits. The two most commonly used techniques for this are the Jordan–Wigner transformation and the Bravyi–Kitaev transformation. However, comparisons of these schemes have previously been limited to individual small molecules. In this paper, we discuss resource implications for the use of the Bravyi–Kitaev mapping scheme, specifically with regard to the number of quantum gates required for simulation. We consider both small systems, which may be simulatable on near-future quantum devices, and systems sufficiently large for classical simulation to be intractable. We use 86 molecular systems to demonstrate that the use of the Bravyi–Kitaev transformation is typically at least approximately as efficient as the canonical Jordan–Wigner transformation and results in substantially reduced gate count estimates when performing limited circuit optimizations.



1. INTRODUCTION

Computational chemistry is the use of well-developed theoretical techniques and algorithms to solve chemical problems. These typically relate to the properties of molecules and chemical reactions. Such processes occur as a result of the rearrangement of electrons among atoms. Quantum chemistry is the branch of computational chemistry concerned with the theoretical understanding of these processes.¹

Although a vast spectrum of methods has been developed for this purpose, the field is restricted by the computational difficulty of the task. Many calculations of interest involve the determination of ground state electronic wave functions and their corresponding energies. To achieve this exactly (to within nonrelativistic assumptions and basis set limitations) requires the *full configuration interaction* approach. This scales factorially with respect to the number of basis functions considered, limiting application of the technique to small molecules.^{2,3}

As this conceptually simple approach is computationally intractable, the difficulty of the task is often reduced by invoking various approximations, well-studied in computational chemistry. While these methods often allow high degrees of precision, some computational tasks—for instance, in determining reaction kinetics or dynamics—would benefit from the decreased error of a full configuration interaction approach.

Quantum simulation algorithms are expected to be capable of alleviating some of the difficulty associated with this through the use of a scalable quantum computer. A quantum computer operates on *qubits*—the quantum equivalent of classical bits.

Instead of a unit which can either have a state value of 0 or 1, qubits exist as superpositions of $|0\rangle$ and $|1\rangle$ states, i.e., $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. A system of n qubits can exist in a superposition of 2^n basis states. Similar to classical computation, operations which manipulate the state of qubits are described as quantum gates and are analogous to classical logic gates. A sequence of quantum gates, intended to perform a computational task, is referred to as a quantum circuit. Gates that perform an operation which entangle the state of two or more qubits are called entangling gates.

Quantum algorithms to address various chemical tasks have been developed, including the determination of energy spectra,⁴ reaction rates,^{5–7} and reaction dynamics.⁸ Quantum simulation of quantum systems—particularly chemistry—is often cited as being one of the most significant potential uses of quantum computation.⁹

The development of a scalable quantum computer is an extremely difficult task. Demonstrations of the quantum simulation of electronic structure problems have mostly thus far been limited to the consideration of small hydrides using a minimal basis set. These have been reported in photonic,¹⁰ NMR,¹¹ and superconducting¹² devices. The first fully scalable demonstration of this kind was performed in 2016.¹² Recently, the use of the variational quantum eigensolver algorithm¹³ has extended this to the simulation of beryllium hydride.¹⁴

Received: May 8, 2018

Published: September 6, 2018

However, recent hardware developments have yielded devices that are rapidly increasing in size.^{12,15} Devices of up to 50 qubits are likely to be available in the near future.¹⁶ It is likely that such devices will bring the field close to the ability to perform classically intractable chemistry calculations.⁷ For this, advances in both hardware and circuit design are necessary.

The canonical quantum algorithm for the solution of the electronic structure problem involves several steps.⁴ First, the molecular orbitals forming a basis for the electronic states must be represented using states of qubits. The electronic Hamiltonian must then be mapped to an operator on the qubit Hilbert space.^{17–20} Following this, the Trotter–Suzuki approximation^{21–23} is invoked to form an evolution operator which is implementable on the quantum device. Finally, a phase estimation algorithm is invoked in order to ascertain the ground state molecular energy.²⁴

Many algorithmic developments have been made to further this goal. In particular, hybrid quantum-classical schemes have been shown to yield accurate results for a fraction of the cost of canonical quantum simulation techniques.^{12–14,25} However, these techniques still require a transformation of electronic states and operators to states of and operations upon qubits. The two most commonly used forms of this transformation are the Jordan–Wigner transformation and the Bravyi–Kitaev transformation,^{4,17–20} although other constructions have been examined.^{26,27}

In the asymptotic limit and without further circuit optimization, the Bravyi–Kitaev transformation is known to have logarithmically superior scaling with respect to circuit length.^{18,19} An examination of the performance of this process requires generating descriptions of quantum circuits which would perform the simulation. Initial assessments of this technique demonstrated a saving associated with the simulation of the hydrogen molecule in a minimal basis, a smallest-case example.¹⁹ This saving was thus expected to also be present for larger chemical examples. However, further investigation of methane revealed that overall gate savings were relatively modest, although substantial savings were yielded in terms of entangling gates.²⁰ To date, no large scale numerical analysis involving many systems has been performed.

In this paper, we use 86 molecular systems to demonstrate that the use of the Bravyi–Kitaev transformation typically results in quantum circuit lengths equal to or shorter than circuits using the canonical Jordan–Wigner transformation. We also consider the impact of Trotter ordering upon both overall gate count and the relative performance of the Bravyi–Kitaev transformation. Varying the Trotter ordering can impact the error incurred in the use of this approximation, potentially resulting in increased overall gate count, even if within each Trotter step the gate count is reduced. As such, we consider the impact of Trotter ordering on the Trotter error, by examining a subset of our systems.

We begin by providing a brief overview of the theory underlying the Bravyi–Kitaev mapping and of Trotterization. In section 3, we discuss circuits with a limited degree of optimization. Following this, we introduce the impact of Trotter ordering, considering its impact on single Trotter step circuit length in section 4, and on the Trotter error in section 5.

2. THEORY

The electronic Hamiltonian in the second quantized formalism is given by

$$\hat{H} = \sum_{ij} h_{ij} a_i^\dagger a_j + \frac{1}{2} \sum_{i,j,k,l} h_{ijkl} a_i^\dagger a_j^\dagger a_k a_l \quad (1)$$

where h_{ij} and h_{ijkl} are Coulombic overlap and exchange integrals determined by the basis set chosen.^{1,28}

Although the number of h_{ij} and h_{ijkl} integrals scales quartically with respect to the number of basis orbitals, they are efficiently computable using conventional, classical computing methods. Additionally, despite this theoretical quartic scaling, the number of nonzero integrals is often substantially reduced through consideration of molecular symmetries.^{1,28} Rather, the difficulty is due to the dimension of the Fock space on which this Hamiltonian acts, which scales exponentially with the number of basis orbitals considered. Restricting the problem to a subspace with a fixed number of electrons reduces this scaling to being factorial with respect to the number of electrons, but for practical problems this remains intractable. This difficulty typically prevents the use of full configuration interaction calculations for purposes other than benchmarking.

On a quantum computer, this scaling difficulty is not present. The task proceeds in four stages, illustrated in Figure 1. First, a

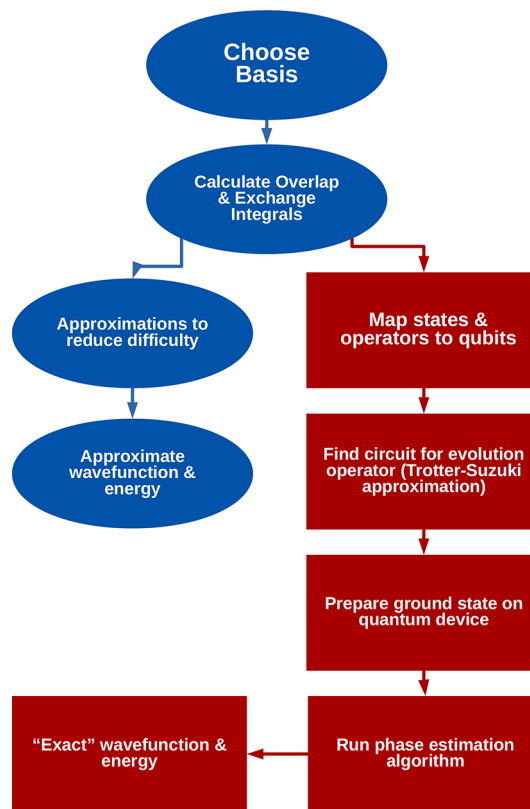


Figure 1. A comparison of classical and quantum algorithms for simulation of electronic structure. Left: classical. Right: quantum.

representation of the Hamiltonian and the molecular orbital spaces upon which it acts must be represented using states of and operations upon qubits. Following this, a Trotter–Suzuki approximation^{21–23,29} is invoked in order to find a quantum circuit to simulate the evolution of the system under the molecular Hamiltonian. The use of this approximation results in an additional error in simulation. However, this error can be arbitrarily reduced through increasing the number of Trotter steps, resulting in only an extra multiplicative factor in the

quantum computational expense. A full consideration of Trotterization error is included in section 5. Having developed such a circuit, a good approximation to the simulated ground state of the system is prepared, and the phase estimation algorithm is used to ascertain the ground state energy of the system.

We initially concern ourselves with the first of these stages—the mapping technique chosen to transform electronic states and operators to states and operators of qubits. We discuss the difference in resource implications for two options for this: the canonical Jordan–Wigner transformation and the Bravyi–Kitaev transformation. In section 5, we address the implications of these mappings when performing a Trotter–Suzuki approximation.

Our task in establishing an appropriate mapping is to find qubit analogues of both the electronic states and the creation and annihilation operators in eq 1. Traditionally, the simplest encoding scheme to determine these is the Jordan–Wigner transformation.^{4,17} Here, n qubits are used to store the occupation number of n electronic spin–orbitals, forming what is known as the occupation basis. If the i th molecular orbital is occupied, then the corresponding i th qubit is in the $|1\rangle$ state, whereas if the molecular orbital is unoccupied, then the qubit is in the $|0\rangle$ state. We then require a representation of the electronic creation and annihilation operators that act on the qubit space, which perform the following set of operations:

$$\begin{aligned}\hat{Q}^+|0\rangle &= |1\rangle & \hat{Q}^+|1\rangle &= 0 \\ \hat{Q}^-|0\rangle &= 0 & \hat{Q}^-|1\rangle &= |0\rangle\end{aligned}\quad (2)$$

A naive assessment would suggest that the standard Pauli σ_i^+ and σ_i^- operators would suffice for this purpose; however, these do not obey the required anticommutation relations:

$$\{a_i^\dagger, a_j^\dagger\} = \{a_i, a_j\} = 0 \quad (3)$$

$$\{a_i^\dagger, a_j\} = \delta_{ij}I \quad (4)$$

For these to hold, the parity of the occupation numbers of the orbitals with index less than i must be calculated, and a phase shift introduced when the parity is odd. This is accomplished by performing a sequence of Pauli Z operations on the preceding qubits, yielding the following:

$$\begin{aligned}a_i^\dagger &= \frac{1}{2}(X_i - iY_i) \otimes_{j<i} Z_j \\ a_i &= \frac{1}{2}(X_i + iY_i) \otimes_{j<i} Z_j\end{aligned}\quad (5)$$

Note that this mapping requires $O(N)$ qubit operations to simulate one electronic operation.

An alternative scheme was envisaged by Bravyi and Kitaev, whereby parity information is stored in the qubit states (i.e., qubit i stores the sum (modulo 2) of the occupation of all electronic states with index less than or equal to i). This basis—the *parity basis*—avoids the additional cost of determining the parity, as this information can be queried with only a single qubit operation.^{18,19} However, this mapping has instead delocalized information regarding the occupation of each electronic orbital. Clearly, any electronic creation or annihilation operation on an orbital with index i requires the update of all qubits with index greater than or equal to i . Consequently, using this mapping the number of qubit operations required to simulate one electronic operation is also $O(N)$.

2.1. Bravyi–Kitaev Mapping. The Bravyi–Kitaev^{18–20} mapping is an attempt to improve upon the linear scaling of the occupation and parity bases. In essence, it is a middle ground between these approaches. For a molecular orbital basis of size N , there are again N qubits used.³⁰ However, the information stored within each qubit now varies dependent on the index i . Note that we begin indexing at $i = 0$. Where i is even, qubit i stores the occupation number of orbital i , as in the Jordan–Wigner mapping. Where i is odd, the qubit stores the parity of a particular set of occupation numbers. When $\log_2(i+1)$ is an integer, the qubit stores the parity of the occupation numbers of all orbitals with indices less than or equal to i . For other cases, the qubit stores the parity of the occupation numbers of orbitals in subdividing binary sets. This complex mapping is best understood through consideration of the matrix which transforms a vector of orbital occupations to qubit states. For example, in the eight orbital/qubit case, this is given by

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} o_0 \\ o_1 \\ o_2 \\ o_3 \\ o_4 \\ o_5 \\ o_6 \\ o_7 \end{bmatrix} = \begin{bmatrix} o_0 \\ o_0 + o_1 \\ o_2 \\ o_0 + o_1 + o_2 + o_3 \\ o_4 \\ o_4 + o_5 \\ o_6 \\ o_0 + o_1 + o_2 + o_3 + o_4 + o_5 + o_6 + o_7 \end{bmatrix} = \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \\ q_4 \\ q_5 \\ q_6 \\ q_7 \end{bmatrix} \quad (6)$$

Here, each o_i value corresponds to the occupation number of the orbital with index i , and the q_i values correspond to the state of the qubit with index i . Where q_i is 0, qubit i is in the $|0\rangle$ state and similarly where q_i is 1, qubit i is in the $|1\rangle$ state. All sums are performed in modulo 2. The matrix on the left is thus the matrix which transforms orbital occupation numbers to qubit states.

Both occupation and parity information is now stored nonlocally. Inspection of eq 6 shows that this information is stored in a number of qubits which grows logarithmically. Thus, any electronic creation or annihilation operation can be simulated in $O(\log_2 n)$ qubit operations. We omit a detailed proof of this here for reasons of brevity. Further details can be found in refs 18–20.

Despite the superior asymptotic scaling of the Bravyi–Kitaev mapping, it is important to consider the increased overhead of its use. Initial implementations noted that the Bravyi–Kitaev mapping is more efficient than the Jordan–Wigner mapping in the simulation of molecular hydrogen in a minimal basis, the smallest possible nontrivial example.¹⁹ It was thus argued that the overhead was not a significant factor, and the superior scaling effectively dominated in all cases. However, further investigation on methane in a minimal basis revealed that this is not the case.²⁰

One purpose of this paper is to find the point at which this asymptotically superior scaling dominates. Examination of the Bravyi–Kitaev creation and annihilation operators permits a rough estimate of this. Note that the qubit creation and annihilation operator equivalents using the Bravyi–Kitaev transformation are given by

$$a_i^\dagger = \frac{1}{2}(X_{U(i)} \otimes X_i \otimes Z_{P(i)} - iX_{U(i)} \otimes X_i \otimes Z_{P(i)}) \quad (7)$$

$$a_i = \frac{1}{2}(X_{U(i)} \otimes X_i \otimes Z_{P(i)} + iX_{U(i)} \otimes X_i \otimes Z_{P(i)}) \quad (8)$$

where $U(i)$ is the “update set” of qubit i , and $P(i)$ is the “parity set” of qubit i . For brevity we do not discuss these sets here; however, their size is of maximum order $O(\log_2 i)$. These expressions are valid only in the case of even i . However, this does not affect our rough estimate. Examining these equations, it is evident that at most $4\log_2 i + 2$ gates are required for the simulation of one Fermionic operation. This quantity is smaller than the simple i gates of the Jordan–Wigner mapping when $i \geq 19$. Thus, noting that the Bravyi–Kitaev mapping is most efficient when N approaches a power of 2 (as it can take increased advantage of its binary tree structure), we conservatively estimate that this point should be at $N \approx 32$. We thus would expect quantum computational cost to be reduced when using the Bravyi–Kitaev transformation for systems with more than $N \approx 32$ spin–orbitals.

2.2. Trotterization and Simulation. In order to perform the phase estimation algorithm to determine the molecular ground-state energy, a quantum circuit simulating the unitary evolution operator $\hat{U} = \exp(-it\hat{H}/\hbar)$ of the molecular Hamiltonian must be found. This is similarly required when utilizing a variational quantum eigensolver algorithm using a coupled-cluster ansatz.^{12–14,31} The qubit form of the electronic Hamiltonian determined through the Bravyi–Kitaev or Jordan–Wigner transformation consists of a weighted sum of strings of Pauli operations. In order to exponentiate this, a Trotter–Suzuki approximation must be invoked.²² The first-order Trotter–Suzuki expansion is

$$e^{-it/\hbar \sum_i \hat{H}_i} \approx \left(\prod_i e^{-it\hat{H}_i/\hbar} \right)^n \quad (9)$$

where H_i are the Hamiltonian subterms (strings of Pauli operations, in our case), and n is the number of Trotter steps. The overall evolution time is now subdivided into n steps. Increasing the number of Trotter steps decreases the error invoked in this procedure. This yields the evolution operator expressed as a product of exponentiated strings of Pauli operations. Standard techniques can be used to transform these into quantum circuits, as shown in Figure 2. The gates within this




Figure 2. Canonical circuit for the simulation of $\exp(-i\frac{\theta}{2}Y_0Z_1Z_3X_4)$.

circuit can be divided into two types: gates that rotate the state of a single qubit and typically more expensive two-qubit entangling gates. These can be implemented sequentially to form a quantum circuit which simulates the entire evolution operator.

The use of a Trotter–Suzuki approximation results in the introduction of error.^{32–34} This error can be reduced by increasing the number of Trotter steps considered. We consider the impact of this error in section 5.

3. BASIC CIRCUITS

Our code was used to assess the serial quantum circuit length corresponding to the simulation of 86 small molecules and atoms. Molecular structures were gathered from the NIST

CCCBDB database³⁵ optimized at the Hartree–Fock level. Most systems used a STO-3G basis; however, larger Pople basis sets were used in 18 trials. Of these, four systems (CH_2^{2+} , HF, LiH, H_2O) using a 3-21G basis set were examined, with the remainder studying H_2 and HeH^+ . Clearly, this choice of basis is insufficient for an accurate solution of the electronic structure problem. When performing a simulation on a real quantum device, a larger basis set would be chosen as in conventional quantum chemistry methods. Fortunately, the error introduced by our choice of basis is fixed and independent of our choice of quantum methodology. As our benchmarking procedure is not directly concerned with the exact energies predicted, our choice of basis set allows for the simulation of a variety of systems with relatively low computational overhead. Details of the systems studied can be found in Table 1 and in the Appendices. Our systems range in size from 2 to 54 spin–orbitals. While containing systems that are classically intractable to simulate, this number is relatively modest in contrast to simulations that would be performed upon a real quantum device;⁷ however, it allows us to maintain relatively low computational expense (approximately 1 week of CPU time for the largest examples).

Hartree–Fock molecular orbitals and their h_{ij} and h_{ijkl} integrals were obtained using the PSI4 electronic structure theory package³⁶ and the FermiLib PSI4 Plugin.³⁷ Our code was then used to generate Jordan–Wigner and Bravyi–Kitaev Hamiltonians. These are stored symbolically as strings of Pauli operations, as in previous work.^{19,20} The Hamiltonian can be stored as blocks of second quantized terms, potentially grouped according to their character, that is, excitation operations, number operations, and so on. Note that due to molecular symmetries (and the symmetries of the integrals present in eq 1), the terms in eq 1 do not necessarily have independent coefficients.

A basis of Hartree–Fock molecular orbitals was used to describe the system when employing the Jordan–Wigner or Bravyi–Kitaev mapping. Much work has been performed in assessing the performance of other basis choices,^{25,33} with localized basis orbitals showing promise in reducing the number of significant terms in the Hamiltonian. However, as this advantage is gained from reduction of the number of significant overlap integrals, there is no obvious reason to believe that the Jordan–Wigner and Bravyi–Kitaev mappings would perform inequivalently in a predictable manner. Preliminary analysis using natural and orthogonalized atomic orbital integrals provided by collaborators³³ did not suggest any consistent dependence of the performance of the Bravyi–Kitaev mapping (versus the Jordan–Wigner mapping) on the choice of basis considered.

As such, in order to reduce the computational cost of our simulations, this degree of freedom was not considered here. A rigorous demonstration of the independence of the performance of the Bravyi–Kitaev mapping on the basis choice could be considered in future work.

Optimization and Trotterization could be performed on the level of second quantized operators. Other works have taken this approach, maintaining Fermionic terms throughout optimization procedures.^{32,38} Instead, our code does not retain the original Fermionic components of the electronic Hamiltonian and reduces the qubit Hamiltonian to a list of strings of Pauli operations, before attempting circuit-level optimization. While fewer assumptions can be made about the structure of the new Hamiltonian (a fact of relevance in section 6), this

Table 1. A breakdown of systems studied. Note that most of the systems involving a non-minimal basis set were H_2 and HeH^+ systems, as specified in the [Appendices](#). Numbers in parentheses indicate the number of systems where Trotter error was considered, as discussed in [section 5](#)

qubits	1–10	11–20	21–30	31–40	41–50	51–60	total
molecules + radicals	1(1)	3(13)	12(3)	4(0)	8(0)	3(0)	41(17)
atoms	10(4)	7(6)	1(0)	1(0)	0(0)	1(0)	20(10)
ions	2(2)	4(3)	0(0)	2(0)	1(0)	1(0)	10(5)
nonminimal basis sets	4(4)	6(5)	5(2)	0(0)	0(0)	0(0)	15(11)
total	17(11)	30(27)	18(5)	7(0)	9(0)	5(0)	86(43)

approach allows for greater flexibility when ordering terms for Trotterization.

From here, our code allows for generation of quantum circuit objects corresponding to the implementation of one Trotter step of the evolution operator of the qubit Hamiltonian (discussed below).

Neglecting any benefits from optimization at the interfaces of Trotter steps (which would save at most $O(N_{\text{trotter}})$ gates), the number of gates necessary for larger numbers of Trotter steps is a simple multiple of the number of gates necessary for one Trotter step. Consequentially, extending our analysis to higher Trotter numbers was not considered necessary for our initial analysis, although this was performed when considering Trotter error in [section 5](#).

A full treatment of the entire phase estimation algorithm (including error correction) was not performed. This is in contrast to other work that has attempted to characterize the resource implications of performing the full procedure.^{7,32} For our analysis of the raw gate counts of Bravyi–Kitaev circuits, this was not considered important for the above reasons. Note that at present, the largest simulations conducted have been of around 45 qubits.^{39–41} These required extensive specialized computing resources. Performing simulations at this level would have made it impossible to perform a large, multisystem survey. In particular, assessing the point at which Bravyi–Kitaev scaling is expected to dominate (around 32 qubits, as discussed above) would have been problematic. Despite this, a full consideration of the phase estimation algorithm could yield useful results in regard to the Trotter error of simulation. For simplicity, we have only considered here the canonical phase estimation procedure, as the expected benefits of the Bravyi–Kitaev mapping will apply in any system which involved the use of such a mapping scheme.

It is important to note that many of the circuits discussed in this paper are substantial in terms of quantum resources required. For around 50 spin-orbitals (and thus 50 logical qubits), the unoptimized circuits consist of around 10^7 gates. It is likely that the implementation of such circuits on a quantum device would require the use of some form of error correcting code. In order to assess resource implications of circuits within a fault tolerant framework, the number of Clifford and non-Clifford gates within the circuit are counted.⁴² While Clifford gates are considered relatively straightforward to implement in a fault-tolerant manner, the resource implications of performing the non-Clifford gates are assessed by counting the number of T ($\pi/4$ phase rotation) gates required to implement them.

While a thorough analysis of the practicality of implementing these circuits on a quantum device would require consideration of this point, our focus here is on assessing the use of the Bravyi–Kitaev mapping. Observing [Figure 2](#), it is evident that only the central rotation gate is a non-Clifford gate. There is one of these gates for every term in the qubit Hamiltonian.

As the number of terms in the Hamiltonian is the same for either mapping scheme, this implies that the number (and type) of non-Clifford gates is the same regardless of the mapping scheme chosen. This is confirmed by numerical analysis in the circuits discussed below. This implies that the choice of Bravyi–Kitaev or Jordan–Wigner mapping does not impact that T count of the circuit and thus does not impact the difficulty of performing error correction. As such, we do not consider this difficulty here.

However, previous studies have shown^{19,20} that the Bravyi–Kitaev mapping results in a reduction in the number of CNOT gates required, independent of those used in constructing a fault-tolerant representation of the central rotation gate. This comes at the expense of an increased number of single qubit Clifford gates required. In other words, these previous examples showed that the Bravyi–Kitaev mapping traded a reduction in two qubit Clifford gates for an increase in one qubit Clifford gates (along with an overall decrease in the total number of Clifford gates). While this is not of huge impact in a fault tolerant framework (as the T count remains the same), experimental devices in the near future are still likely to benefit from the minimization of entangling gates so as to reduce error. As such, we have considered the breakdown of circuits into entangling and single qubit gates in this paper.

Initially, the Hamiltonians were ordered by the magnitude of their coefficients. This somewhat arbitrary ordering was chosen in order to assess the preoptimization efficacy of the Bravyi–Kitaev mapping and is in contrast to optimized ordering schemes used in other work.^{19,38} These are considered in detail in [sections 4 and 5](#). The systems studied involved between 2 (the hydrogen atom) and 56 (the iodine atom) spin-orbitals. As to be expected,³² the serial circuit length dramatically increases for larger systems, requiring of order 10^7 gates for the simulation of systems involving around 50 spin-orbitals. While not as ruinous as the factorial difficulty of classical full configuration interaction, this large circuit length illustrates the need for compiler optimizations.

An initial assessment of circuits for the implementation of Jordan–Wigner and Bravyi–Kitaev Hamiltonians suggests that the use of the Bravyi–Kitaev mapping is associated with a substantial improvement for the larger systems. [Figure 3](#) demonstrates this. From roughly 30 spin-orbitals, this improvement is consistent and constitutes approximately 25% of the overall circuit length for the largest of the systems we have examined. However, many systems smaller than this see no improvement, or even demonstrate larger circuit lengths. This is in line with our earlier rough estimate that the superior scaling of the Bravyi–Kitaev mapping dominates the increased overhead at around 32 spin-orbitals. Classical full configuration interaction calculations have been performed on systems marginally larger than this (36 molecular orbitals).⁴³ Consequently, for simulations aiming to achieve results which are classically

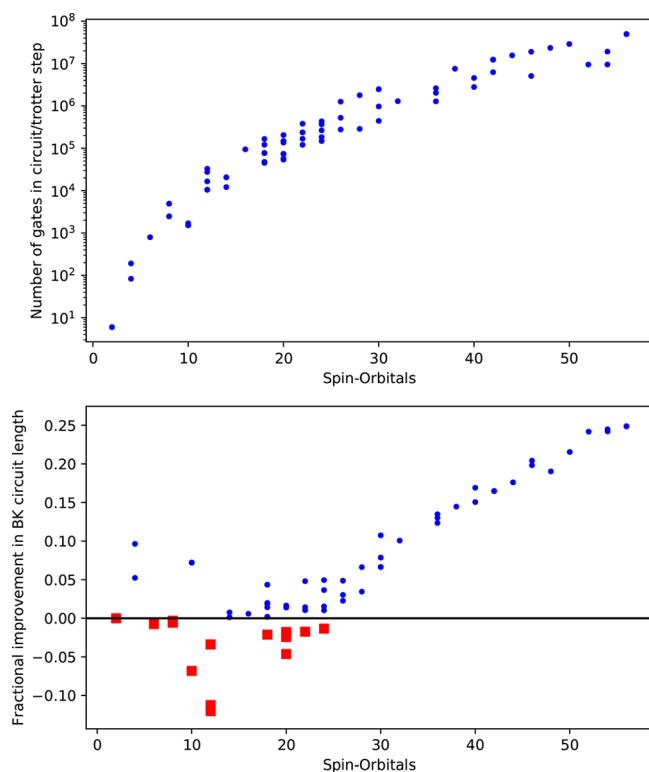


Figure 3. Unoptimized gate counts. Upper: Total number of gates in Jordan–Wigner circuits, before optimization. Lower: Gate savings through use of Bravyi–Kitaev mapping as a fraction of Jordan–Wigner gate count. Squares indicate instances where the Jordan–Wigner transformation outperformed the Bravyi–Kitaev transformation. In this scheme, the Bravyi–Kitaev transformation reliably results in shorter circuits from around 30 spin–orbitals, with up to around 25% gate savings in the examples with 50 spin–orbitals.

intractable, a naive approach involving no circuit optimizations would be substantially eased through the use of the Bravyi–Kitaev mapping.

In addition to optimization performed through combination of duplicate Pauli strings, our code optimizes circuits by the cancellation of duplicate gates.³⁸ Circuit objects can automatically search their gate sequence for duplicate self-inverse gates and remove them. Furthermore, the code tests individual gates for commutativity with gates that follow in sequence. If such commutativity is present, it tests to see if any accessible gates are accessible through commutation. This is performed according to a set of rules: gates acting on different qubits always commute, CNOT gates commute unless one targets the other’s control, and so on. This avoids the generation of matrix representations of the gates. Optimisation in this form is repeated until the circuit is self-consistent and no further optimization could be yielded.⁴⁴

Based on work by Hastings, Wecker, Bauer, and Troyer,^{32,38} savings from this procedure arise from two factors. First, redundancy in parity determination is eliminated, as this information is not recomputed after every term in the Hamiltonian. This results in savings in the CNOT string used to determine parities. Additionally, basis changes are not recomputed when unnecessary. This saves on the single qubit H and Y gates necessary to set these bases.

When duplicate gates in the circuit are removed, the superiority of the Bravyi–Kitaev mapping is even more pronounced while still using a magnitude ordering. This relative

improvement appears to increase with larger circuits, as demonstrated by Figure 4. In circuits involving more than




Figure 4. Number of gates in Bravyi–Kitaev and Jordan–Wigner circuits, before and after gate cancellation, versus the number of gates in Jordan–Wigner circuit before optimization, using a magnitude ordering. Upper: Total gate count. Lower: Entangling gate count only.

about 10^7 gates, the reduction in gate count associated with the use of the Bravyi–Kitaev mapping is larger than that of gate cancellation. In these cases, the use of the Bravyi–Kitaev mapping results in circuits that are approximately 30–40% shorter. Additionally, the number of gates canceled using the Bravyi–Kitaev mapping is several times greater than the number of gates canceled using Jordan–Wigner mapping. In some cases the advantage associated with the Bravyi–Kitaev mapping reduces the circuit length to that observed for systems involving fewer orbitals. For example, the Bravyi–Kitaev circuit for the simulation of the iodine atom (54 spin–orbitals) requires 5 204 912 gates per Trotter step, whereas the Jordan–Wigner circuit for the simulation of acetone (52 spin–orbitals) requires 8 954 933 gates per Trotter step.

With systems requiring more than 10^6 gates to simulate, there are no examples where the optimized Jordan–Wigner technique outperforms the optimized Bravyi–Kitaev technique. Thus, using this magnitude ordering, it is clear that the Bravyi–Kitaev mapping should be preferred to the Jordan–Wigner mapping in the general case.

As discussed above, previous work²⁰ on the methane molecule indicated that the Bravyi–Kitaev mapping may be advantageous with particular regard to the number of entangling gates required. Our findings here confirm that this advantage holds in general, as shown in Figure 4. In addition to the general gate savings associated with the Bravyi–Kitaev mapping, we observe a substantial decrease in the number of CNOT gates

required. We consider this to be a major advantage of the Bravyi–Kitaev mapping. This advantage is typically offset by a small increase in the number of single qubit gates required (as the total savings are smaller than the entangling gate savings).

While using a magnitude ordering, gate cancellation does not result in a great deal of entangling gate savings, which are typically far fewer than the advantage associated with the use of the Bravyi–Kitaev transformation. Instead, the bulk of gate savings associated with cancellation techniques arises from maintaining the calculation basis between sequential terms, as opposed to continually resetting to the computational basis. It is likely that many CNOT strings are being “trapped” behind noncommuting gates earlier in their respective CNOT strings. This could be alleviated by further circuit optimization; however, this task is difficult to perform without further decreasing the locality of the CNOT chain.

The results of this scheme use a magnitude ordering for both the Jordan–Wigner mapping and the Bravyi–Kitaev mapping. Further analysis of the performance of the Bravyi–Kitaev mapping is impossible without consideration of the Trotterization ordering chosen. This is considered in the following section. Manipulation of the Trotter ordering can also cause variation in the Trotter error, which could result in an increased number of Trotter steps necessary for constant precision. This increased difficulty could undermine the savings gained from the use of a particular ordering in terms of cancellation, and is examined in section 5.

4. IMPACT OF TROTTER ORDERING

As discussed above, the overall goal is to find a minimal length circuit that can implement the unitary evolution operator of the quantum Hamiltonian. As no standard circuit for the simulation of the exponentiated total Hamiltonian exists, a Trotter–Suzuki approximation must be invoked (eq 9).

The ordering of terms in this approximation is important. It has been demonstrated³³ that the error due to the utilization of Trotter–Suzuki formulas strongly depends on the term ordering chosen. Additionally, the number of duplicate gates depends strongly on the ordering chosen. Both of these factors influence the length of the overall quantum circuit. The previous sections utilized a magnitude ordering of Trotter terms. This ordering is significantly physically meaningful, as terms with higher magnitude are likely to correspond to stronger physical interactions. However, it is also known to be suboptimal in terms of gate cancellation procedures.³⁸

As the number of potential orderings grows factorially, the problem of finding an optimal ordering scheme is a difficult one. However, ordering schemes that are superior for the process of gate cancellation can be found, as the similarity of sequential Pauli strings—and thus the savings from cancellation—can be determined when specifying the Hamiltonian. Our analysis compares the impact of the use of the Bravyi–Kitaev mapping with a varying choice of ordering.

A lexicographic ordering is expected to maximize the gate savings obtained through cancellation, as the similarity of adjacent terms is maximized.³⁸ As our code optimizes on the level of Pauli strings rather than Fermionic operators, we order on this level also, with no ordering performed on the Fermionic level. Our code stores Pauli strings as lists of base 4 integers. A lexicographic ordering in this scheme is then essentially a bitwise numerical ordering.

We present results based on this ordering explicitly in Figure 5. Note first the dramatic gate savings associated with using this




Figure 5. Number of gates in Bravyi–Kitaev and Jordan–Wigner circuits, before and after gate cancellation, versus the number of gates in Jordan–Wigner circuit before optimization, using a lexicographic ordering. Upper: Total gate count. Lower: Entangling gate count only.

optimization and ordering scheme. Whereas using a magnitude ordering, the Bravyi–Kitaev mapping provided the bulk of the gate savings once gate cancellation had been performed, now the impact of the Bravyi–Kitaev mapping is relatively minor. The savings associated with the use of a lexicographic ordering far outweigh the savings associated with the use of the Bravyi–Kitaev mapping with a magnitude ordering: in the longer circuits included in our analysis, the former are approximately thrice that of the latter. In these circumstances, the Jordan–Wigner and Bravyi–Kitaev mappings appear roughly equivalent for the smaller circuits. The Jordan–Wigner mapping outperforms the Bravyi–Kitaev mapping in some of the longer circuits, considering both total and entangling gate counts. In the longest circuits considered (propanol), the use of the Jordan–Wigner mapping resulted in circuits that were approximately 25% shorter. This is attributed to the relative complexity of the Bravyi–Kitaev mapping resulting in a reduction of linearity in the CNOT chains, which hampers gate cancellation.

The error implications of Trotter ordering schemes are difficult to predict. Bounds exist on the error yielded from Trotterization,³² although these are often loose.³³ A qualitative estimate can be obtained through determination of the norm of the Trotter error operator; however, the quantitative behavior of this still often overestimates the error in real chemical examples.³³ It is useful to compare the implications of the Bravyi–Kitaev mapping when using several ordering schemes. To this end, we repeated our calculations using four ordering schemes. In addition to a single randomized ordering, a lexicographic ordering and an ordering of terms by magnitude, we include an ordering generated by regularly

interspersing terms from the lexicographic and magnitude orderings. Note that this ordering is relatively arbitrary and is intended for comparison purposes. An exhaustive search of the ordering space is clearly intractable for nontrivial systems, owing to the factorial growth of the number of possible orderings.

Nonetheless, our findings are remarkably consistent, with the Bravyi–Kitaev mapping outperforming the Jordan–Wigner mapping in all cases apart from the lexicographic ordering. This advantage increases with the number of spin–orbitals used. Beyond $N = 10$ in all nonlexicographic cases, the advantage associated with the Bravyi–Kitaev mapping exceeds that of just using gate cancellation. This advantage is increased when considering CNOT count and can be dramatic—when using a magnitude ordering, the savings associated with the Bravyi–Kitaev mapping are approximately an order of magnitude greater than those obtained by using gate cancellation alone (Figure 6)

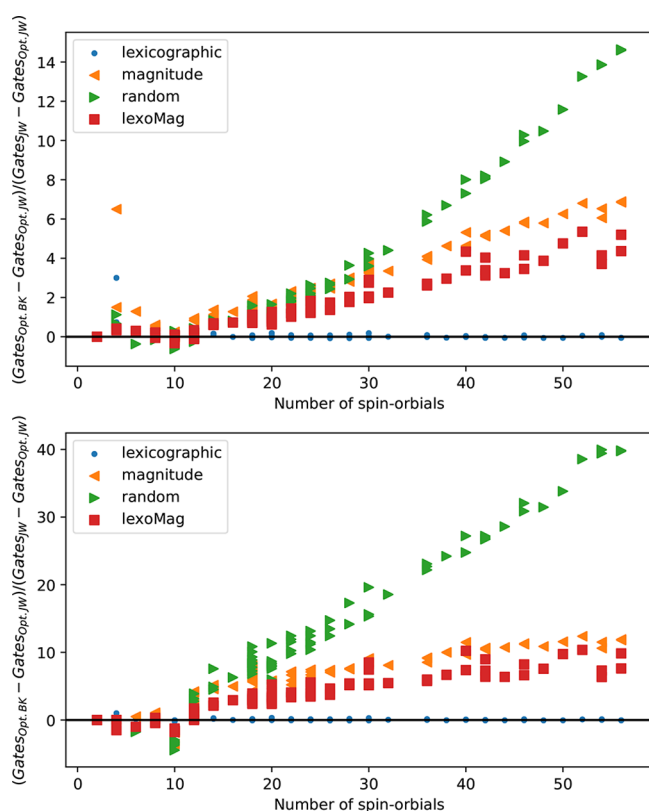


Figure 6. Gate savings associated with the Bravyi–Kitaev mapping normalized by the gate savings acquired using the same optimization with the Jordan–Wigner mapping versus the number of spin–orbitals simulated, for various ordering schemes. A value of 0 indicates that the optimized Bravyi–Kitaev and optimized Jordan–Wigner circuits have equal number of gates. A value of 1 indicates that the optimized Bravyi–Kitaev mapping outperforms the optimized Jordan–Wigner mapping by a number of gates equal to that saved by performing optimization on the raw Jordan–Wigner circuit. Upper: total gates. Lower: entangling gates.

For large systems, searching even a statistically significant subset of the space of possible orderings is clearly intractable, owing to the factorial growth of the number of possibilities. For each system, our random ordering is only one of these myriad choices. Consequentially, it does not represent a statistically meaningful representation of the bulk of the possible

orderings. Nonetheless, it is interesting that this random choice qualitatively manifests the same trend as our other ordering schemes to a very large extent. Quantitatively, the advantage associated with the Bravyi–Kitaev mapping when using these random orderings is dramatic; for our largest example, the reduction in CNOT count using the Bravyi–Kitaev mapping was 40 times the reduction using gate cancellation alone. In our results, the use of a tailored ordering scheme (whether it be ordering lexicographically, by magnitude or otherwise) results in a reduction of advantage for the Bravyi–Kitaev mapping.

Although these figures suggest that the use of the Bravyi–Kitaev mapping results in shorter circuits when using most possible orderings, it should be emphasized that the lexicographic ordering dramatically decreases gate count through cancellation irrespective of mapping strategy. We restate that using this strategy, Figure 5 shows that employing the Jordan–Wigner mapping results in marginally shorter circuits than the Bravyi–Kitaev mapping.

The choice of Trotter ordering may be dictated by other factors - for instance, architecture constraints. In these circumstances, calculations will be dramatically shortened by the use of the Bravyi–Kitaev mapping, as this mapping results in reduced circuit length for all of the nonlexicographic orderings considered.

The choice of ordering has previously been shown to hold significant impact on the Trotter error.^{19,32} Potentially, minimization of Trotter error may therefore require an ordering being chosen which is suboptimal in terms of single Trotter step gate count, where the Bravyi–Kitaev holds a significant advantage.

As such, consideration of the Bravyi–Kitaev mapping within the context of different ordering schemes requires an estimation of the associated Trotter error. Previous work^{19,20} indicates that for the hydrogen and methane molecules, the use of the Bravyi–Kitaev mapping can result in a reduced Trotter error—although, in the latter case, insufficiently to decrease the number of Trotter steps required for accurate simulation. We consider these points in the general case in the following section.

In general, we conclude here that the use of the Bravyi–Kitaev mapping does not result in a predictable improvement in gate count when using an ordering optimized for gate cancellation, ignoring Trotter error implications. This is in contrast with the substantial and predictable improvement observed with other orderings, particularly when a random ordering is used.

5. TROTTER ERROR CONSIDERATIONS

As discussed above, the use of Trotter–Suzuki approximations cause error which decreases with the number of Trotter steps performed. Bounds on and estimates of this error have been established;³² however, it has been shown³³ that these estimates often overestimate the actual error incurred by many orders of magnitude. Exact determination of the Trotter error is exponentially hard, as the exact ground state energy must be determined in advance to serve as a reference.

Having generated Jordan–Wigner and Bravyi–Kitaev Hamiltonians as symbolic lists of Pauli strings, our code can proceed in several ways. For smaller systems, a sparse matrix representation of these Hamiltonians can be generated using SciPy's⁴⁵ sparse matrix methods. This can be diagonalized to find an exact ground-state eigenvalue (to compare against the estimate provided by further code) and a ground-state eigenvector. These can be compared against traditional full

configuration interaction calculations obtained through direct diagonalization of the Hamiltonian for verification purposes.

Our code can also generate Trotterized Jordan–Wigner and Bravyi–Kitaev Hamiltonians while maintaining the symbolic representation. The action of these Hamiltonians on a given state can be simulated. Performing this on the generated ground state eigenvectors allows for the determination of Trotter error without the storage difficulty of repeatedly generating the Trotterized evolution operator in matrix form. Nonetheless, the initial determination of ground state eigenvectors remains exponentially difficult, requiring on the order of hundreds of gigabytes of RAM for examples with more than 20 spin–orbitals. Consequently, we restricted our error analysis to these smaller systems, as indicated by Table 1 and the Appendices.

We conducted this analysis for 34 of our previously discussed systems. The calculations were performed for a variety of choices of Trotter order and step number. Nonetheless, it is remarkable that limiting our discussion to trials of one Trotter step of first order suffices to yield chemical accuracy (i.e., to within 1 kcal/mol of the FCI energy) for a single application of the evolution operator. For simulation of the full-phase estimation procedure, determination of the error incurred in the application of higher powers of this operator would be necessary. However, this does not qualitatively effect our ordering comparison.

Figure 7 demonstrates the results of these simulations. Encouragingly, in the overwhelming majority of systems considered, the Trotter error is extremely small even with only one Trotter step: it is often less than 0.001 hartree. It is possible that this is an artifact of the small number of spin orbitals in the systems considered in our error analysis. It is also worth noting that these errors will be compounded when considering higher bits of precision in a full phase estimation procedure. Nonetheless, it does suggest that the number of Trotter steps required for chemically accurate simulation of larger systems will be relatively modest, potentially less than 10 Trotter steps for the first bit of precision.

Figure 7 additionally demonstrates the ordering dependence of the Trotter error. We consider the Trotter error of each systematic (i.e., nonrandom) ordering normalized by the Trotter error of the lexicographic ordering for each system.

In most cases, the magnitude ordering appears to achieve a lower Trotter error than the lexicographic ordering. In some cases, this difference exceeds an order of magnitude. However, it is important to note that this represents a large variance on an exceptionally small error. Noting that one Trotter step was sufficient for chemical accuracy in most of the systems studied here, we do not argue that this indicates that a magnitude ordering achieves a useful reduction in error compared to a lexicographic ordering. Future work is required to investigate how significant this distinction is when propagated through the entire phase estimation procedure, as in these circumstances this effect could become sufficiently significant to determine ordering choice.

An examination of the relative performance of the Bravyi–Kitaev mapping and the Jordan–Wigner mapping in the context of ordering strategies is included as Figure 8. Again, the distinction between the two mappings is in most cases not as substantial as the differences observed for single Trotter step circuit length. In the majority of systems, the normalized difference between the two errors is between 1 and -1 —that is, the error associated with one mapping is very rarely more than double that of the other. Several systems display substantially

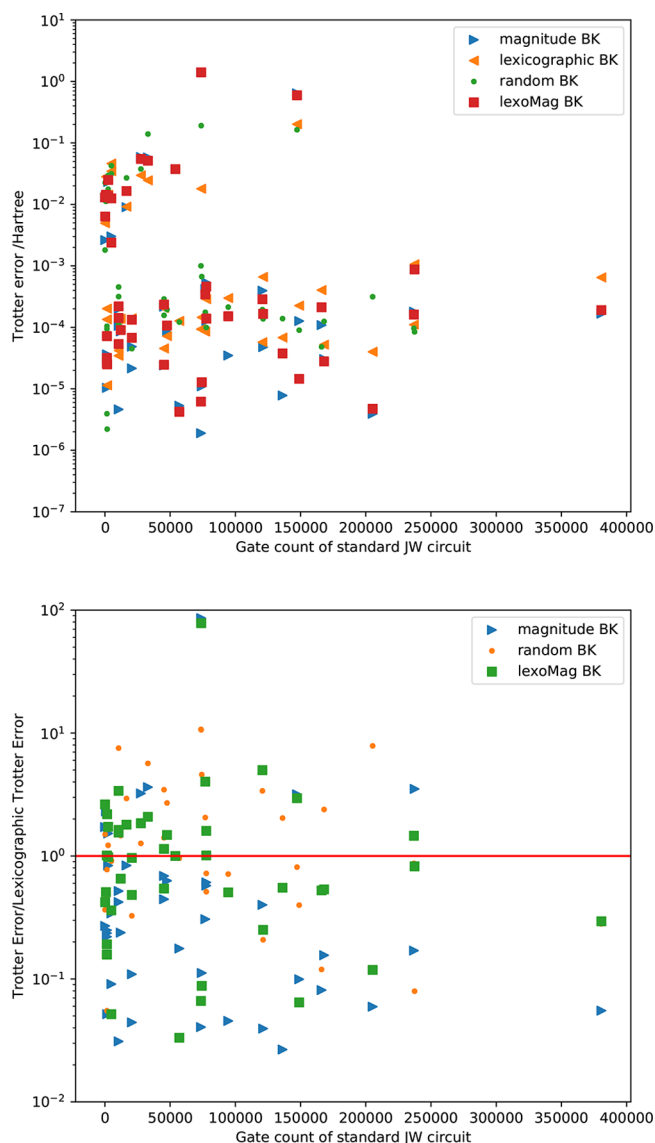


Figure 7. Trotter error using the Bravyi–Kitaev mapping as a function of Jordan–Wigner circuit length, for different ordering schemes. Upper: Absolute Trotter error. Middle: Trotter error relative to error of lexicographic ordering. Note this excludes two instances where the relative Trotter error was >20 . The red line indicates a relative performance of 1 (i.e., below the line, the ordering results in a lower Trotter error than a lexicographic ordering). Note that the magnitude ordering usually results in a substantially lower error; however, in most of these cases, the Trotter error was already around 10^{-4} .

increased error associated with the Bravyi–Kitaev mapping (including two not shown on Figure 8 for scale clarity); however, in the general case no such pattern emerges. Using a lexicographic ordering, a preference for the Bravyi–Kitaev mapping is observable in most systems. In some cases, the error is almost halved by the use of the Bravyi–Kitaev mapping. At larger system sizes this could become a more substantial effect; however, we do not contend that our data provides concrete evidence as to whether this is true. For a magnitude ordering the Jordan–Wigner and Bravyi–Kitaev mappings yield nearly identical Trotter error in almost all cases. In a sense, this could be attributed to the more directly physical nature of the magnitude ordering. Important terms will intrinsically be simulated earlier in sequence using both mappings.

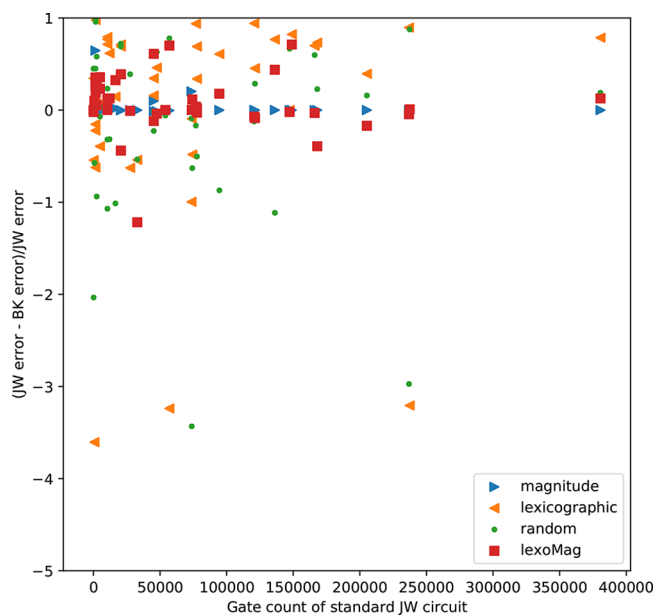


Figure 8. Relative Trotter error of Jordan–Wigner and Bravyi–Kitaev mappings as a function of gate count. The difference in error is normalized by the Jordan–Wigner error, such that a value of 0 indicates equivalent performance, with negative values implying high Bravyi–Kitaev error. Both schemes display remarkably similar errors for the magnitude ordering, likely due to the high degree of “physicality” of the ordering. For a lexicographic ordering, the Bravyi–Kitaev mapping shows lower error in most of the systems studied.

As such, the error of both is likely to be similar in this case. As to be expected, the “lexoMag” ordering performs roughly as a combination of the magnitude and lexicographic orderings.

As above, the impact of this variation in error could prove substantial when propagated through the entire phase estimation algorithm. Using a lexicographic ordering—optimal for gate cancellation—the Bravyi–Kitaev mapping outperforms the Jordan–Wigner mapping in most cases. If this superior error performance scales to above 30 qubit systems, this could result in a reduction of the necessary Trotter steps for simulation. The consequent reduction in circuit length could counterbalance the marginally increased individual Trotter step cost of using the Bravyi–Kitaev mapping in a lexicographic ordering. Further work examining the entire procedure should assess this. Nonetheless, it should not be forgotten that an exact determination of the Trotterization error is equivalent to the solution of the exponentially hard eigenvalue problem itself. Consequentially, this approach may prove to be intractable. Qualitatively, an examination of the norm of the Trotter error operator may prove to be instructive.³³

6. FURTHER OPTIMIZED CIRCUITS

Our code also generates some of the optimized circuits developed by Hastings, Wecker, Bauer, and Troyer,³⁸ in order to assess the impact of these optimizations with respect to the Bravyi–Kitaev mapping. Examining Figure 2, it is evident that many of the CNOT strings may be “blocked off” from cancellation by the basis change gates exterior to them. In this approach, these basis change gates are brought inside the bulk of the CNOT string, as shown in Figure 9. Note that in the first CNOT string, the final CNOT is replaced by a CZ gate. An inspection of the gate sequence implemented on the final qubit in the case of even and odd parities of each subset of

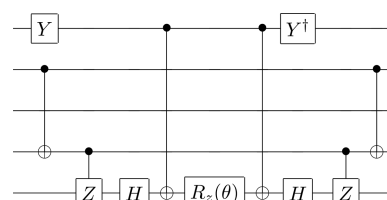


Figure 9. A circuit performing an equivalent operation to Figure 2, using a basis change shift optimization.

qubits demonstrates why this is the case, as discussed in the Appendices.

Our implementation is a slight modification of this technique. As discussed above, our code reduces the Hamiltonian to a symbolic list of exponentiated Pauli strings which does not preserve the electronic Hamiltonian’s original H_{pqrs} components. In this scheme, it is not always the case that the “final” qubit—the qubit which is acted on by the central single qubit rotation—is always in the X basis. As such, the above method requires modification. For example, if the central qubit is to be rotated in the Z basis, the additional CZ gate could simply be commuted through the central rotation and canceled, leading to a phase error.

Fortunately, in the case that the central qubit is to be rotated in either the Y or Z basis, a CNOT gate can be used in place of the CZ gate, as in Figure 9. To demonstrate this, we

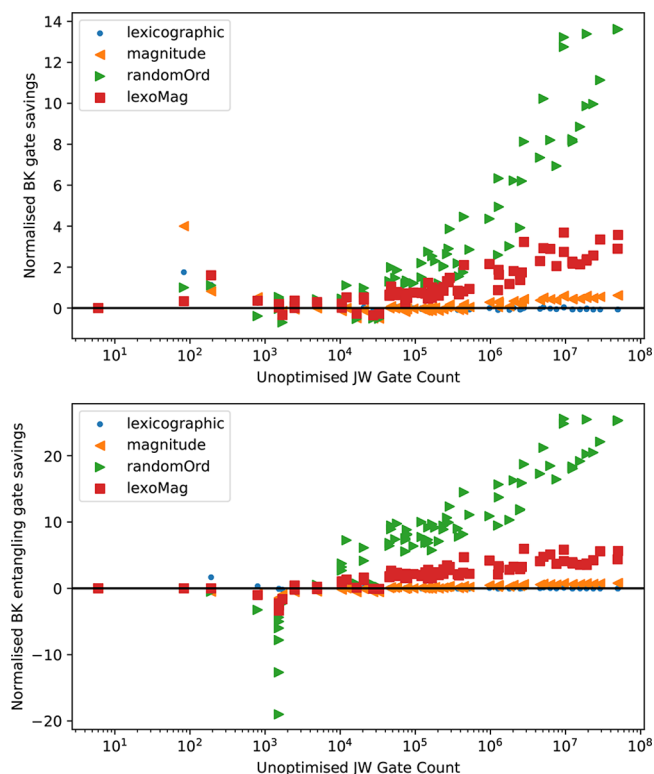


Figure 10. Gate savings associated with the Bravyi–Kitaev mapping normalized by the gate savings acquired using the same optimization with the Jordan–Wigner mapping, using the modified basis change shift technique. Upper: Total gate count. Lower: Entangling gate count. Here we see that the Bravyi–Kitaev and Jordan–Wigner mappings perform relatively equivalently when using both lexicographic and magnitude orderings. Using an alternative ordering, the Bravyi–Kitaev mapping is superior, however these result in an overall increased gate count in both cases.

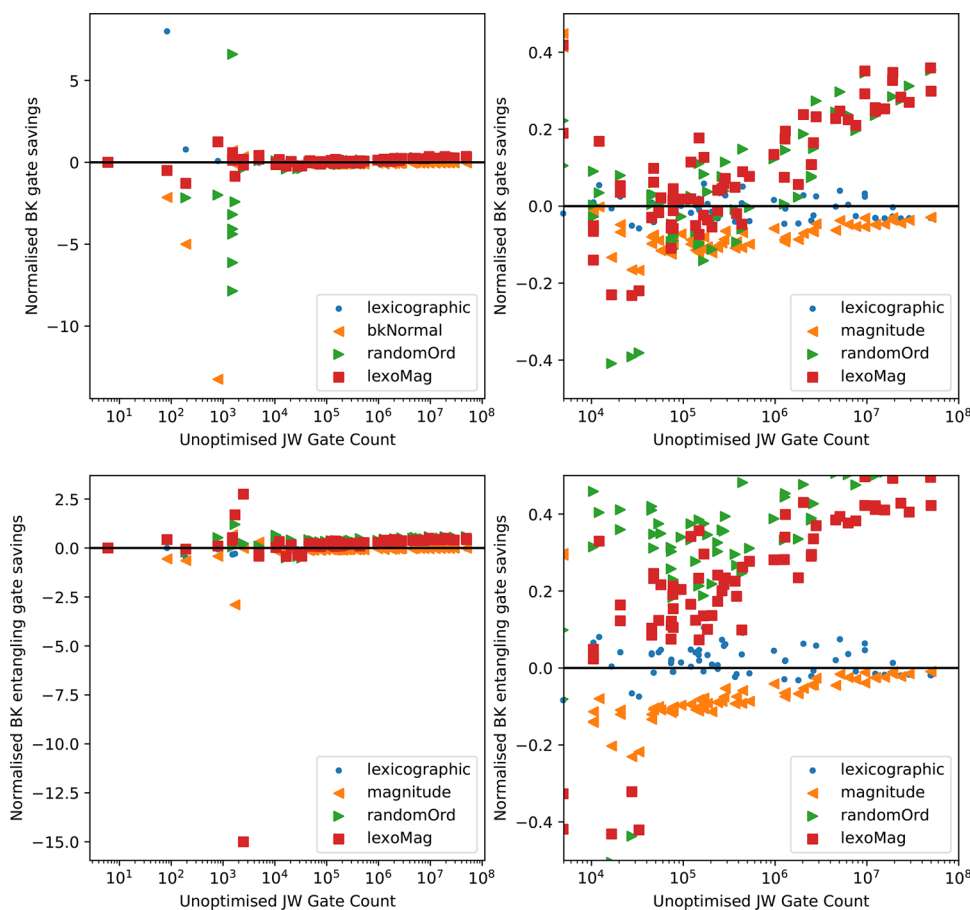


Figure 11. Gate savings associated with the Bravyi–Kitaev mapping normalized by the gate savings acquired using the same optimization with the Jordan–Wigner mapping, using the ancilla-based technique. Upper left: Total gate count. Upper right: Total gate count, zoomed. Lower left: Entangling gate count. Lower right: Entangling gate count, zoomed. Here, we again see roughly equivalent performance when using either a magnitude or lexicographic ordering. Once again, when using an alternative ordering, the Bravyi–Kitaev mapping is superior, however to a greatly reduced extent.

consider the action on the final qubit in the case of the parity of each subset of qubits, as shown in the [Appendices](#).

Additionally, circuits described by Hastings, Wecker, Bauer, and Troyer³⁸ involving an ancilla qubit can be generated. Here, all parity calculating CNOT gates are targeted on a single ancilla qubit. This allows for CNOTS performed in the same basis to be moved around arbitrarily, allowing for a greater level of gate cancellation.

Circuits of these forms are known to reduce overall gate count substantially.³⁸ We focus here on the relevance of the Bravyi–Kitaev mapping when using these techniques.

Using the former technique, the performance of the Bravyi–Kitaev technique using a lexicographic or random ordering displays roughly the same trend as previous circuits. However, using a magnitude ordering reduces the efficacy of the Bravyi–Kitaev mapping to the point of near-equivalence to the Jordan–Wigner mapping. Any savings or penalties associated with the use of the Bravyi–Kitaev mapping are then negligible compared to gains from the basis change shift optimization procedure. We do not yet have an explanation for the ordering dependence of this behavior. Nonetheless, the Bravyi–Kitaev mapping never substantially under-performs when compared to the Jordan–Wigner mapping. In essence, the observed trends are the same as for the previous circuits, albeit with a greatly reduced factor of improvement associated with the use of the Bravyi–Kitaev mapping ([Figure 10](#)).

The use of ancilla circuits displays a markedly different trend. Here, the effect of the Bravyi–Kitaev mapping is greatly reduced in all ordering schemes—there is a maximum of around 30% reduction in the largest examples, when using a “lexoMag” ordering. It is likely that any major savings or penalties associated with the Bravyi–Kitaev mapping are masked by the substantial savings associated with the use of ancillarised circuits. Using a lexicographic ordering, there is no predictable difference between the two mapping schemes at all.

Curiously, using a magnitude ordering reverses the trend observed for previous optimization schemes. Here, the Bravyi–Kitaev mapping is consistently outperformed by the Jordan–Wigner mapping. However, this distinction is relatively small, and disappears entirely in larger system sizes. As such, we do not conclude that a consistent preference for either mapping is present using these circuits. Note that the use of such circuits may be undesirable in certain architectures, due to the loss of nearest-neighbor CNOT chains, which could undermine the substantial savings associated with this technique ([Figure 11](#)).

Finally, we note that in all optimization systems studied, a random choice of ordering resulted in the strongest performance of the Bravyi–Kitaev mapping. While one ordering is clearly not a statistically meaningful sample of the possible orderings, it is interesting that our results here are extremely consistent.

Table 2. All Molecules Considered in This Study^a

name	charge	mult.	basis	no. qubits	level	name	charge	mult.	basis	no. qubits	level
methane	0	1	STO-3G	18	Y	He	0	1	STO-3G	2	N
ethane	0	1	STO-3G	32	N	Li	0	2	STO-3G	10	Y
ethene	0	1	STO-3G	28	N	Be	0	1	STO-3G	10	Y
propene	0	1	STO-3G	42	N	B	0	2	STO-3G	10	Y
ethyne	0	1	STO-3G	24	N	C	0	3	STO-3G	10	Y
methanol	0	1	STO-3G	28	N	N	0	4	STO-3G	10	N
ethanol	0	1	STO-3G	42	N	O	0	3	STO-3G	10	N
isopropanol	0	1	STO-3G	56	N	F	0	2	STO-3G	10	N
1-propanol	0	1	STO-3G	56	N	Ne	0	1	STO-3G	10	N
propanone	0	1	STO-3G	52	N	Na	0	2	STO-3G	18	Y
methanal	0	1	STO-3G	24	N	Mg	0	1	STO-3G	18	Y
ethanal	0	1	STO-3G	38	N	Si	0	3	STO-3G	18	Y
ethanoate ion	-1	1	STO-3G	46	N	P	0	4	STO-3G	18	Y
ethanoic acid	0	1	STO-3G	48	N	S	0	3	STO-3G	18	Y
hydrogen peroxide	0	1	STO-3G	24	N	Cl	0	2	STO-3G	18	Y
ethanamide	0	1	STO-3G	50	N	Ar	0	1	STO-3G	18	N
methylamine	0	1	STO-3G	30	N	K	0	2	STO-3G	26	N
dimethylamine	0	1	STO-3G	44	N	I	0	2	STO-3G	54	N
ammonia	0	1	STO-3G	16	Y	Br	0	2	STO-3G	36	N
ammonium	1	1	STO-3G	18	Y	I ⁻	-1	1	STO-3G	54	N
nitrogen dioxide	0	2	STO-3G	30	N	Br ⁻	-1	1	STO-3G	36	N
lithium hydroxide	0	1	STO-3G	22	Y	Cl ⁻	-1	1	STO-3G	18	N
sodium hydroxide	0	1	STO-3G	30	N	H ₂	0	1	3-21G	8	Y
H ₂	0	1	STO-3G	4	Y	H ₂	0	1	6-31G	8	Y
lithium hydride	0	1	STO-3G	12	Y	H ₂	0	1	6-31G**	20	Y
beryllium hydride	0	1	STO-3G	14	Y	H ₂	0	1	6-311G*	12	Y
N ₂	0	1	STO-3G	20	Y	H ₂	0	1	6-311G**	24	N
O ₂	0	3	STO-3G	20	Y	HeH ⁺	1	1	STO-3G	4	Y
O ₂	0	1	STO-3G	20	Y	HeH ⁺	1	1	3-21G	8	Y
F ₂	0	1	STO-3G	20	Y	HeH ⁺	1	1	6-31G	8	Y
sodium hydride	0	1	STO-3G	20	Y	HeH ⁺	1	1	6-31G**	20	Y
magnesium hydride	0	1	STO-3G	22	Y	HeH ⁺	1	1	6-311G*	12	Y
Cl ₂	0	1	STO-3G	36	N	HeH ⁺	1	1	6-311G**	24	N
HCl	0	1	STO-3G	20	Y	methylene	0	3	3-21G	14	N
HF	0	1	STO-3G	12	Y	HF	0	1	3-21G	22	Y
carbon dioxide	0	1	STO-3G	30	N	lithium hydride	0	1	3-21G	22	Y
carbon monoxide	0	1	STO-3G	20	Y	H ₂ O	0	1	3-21G	26	N
water	0	1	STO-3G	14	Y	H ₃ ⁺	1	1	STO-3G	6	Y
methylene	0	3	STO-3G	14	Y	H ₃ ⁺	1	1	3-21G	12	Y
hydroxide	-1	1	STO-3G	12	Y	CO ₃	0	1	STO-3G	40	N
HNO ₃	0	1	STO-3G	42	N	magnesium hydroxide	0	1	STO-3G	42	N
nitrate	-1	1	STO-3G	40	N	H ₂ S	0	1	STO-3G	22	Y
H	0	2	STO-3G	2	N						

^aLevel corresponds to whether or not the error incurred in Trotterization was considered: N indicates that this analysis was not performed for this system, Y indicates that it was.

7. CONCLUSIONS

In this paper, we have made a detailed comparison of the Jordan–Wigner and Bravyi–Kitaev mappings using a variety of advanced circuit optimization techniques drawn from the theory of quantum simulation. Using unoptimized circuits, the use of the Bravyi–Kitaev mapping dramatically reduces the quantum computational expense of simulation in all systems involving more than 30 qubits and for systems likely to be classically intractable to simulate. In cases with approximately 50 qubits, this improvement reduced the gate count by roughly 25%.

The use of an optimized Trotter ordering absorbs the advantage associated with the Bravyi–Kitaev mapping. The Jordan–Wigner mapping in this case results in slightly shorter circuits for an individual Trotter step. Nonetheless, in most

cases the use of the Bravyi–Kitaev mapping is at worst roughly equivalent to the use of the traditional Jordan–Wigner mapping. Frequently, the gate count reduction is particularly large in the number of expensive entangling gates required. Notably, the Bravyi–Kitaev mapping is superior in all cases aside from the lexicographic ordering.

While our results suggest a slightly reduced error associated with the use of a magnitude ordering, we do not conclude that this ordering should be favored due to the substantial overall gate count associated with a lexicographic ordering. Our analysis of Trotter ordering error suggests that the use of the Bravyi–Kitaev mapping typically results in a reduced Trotter error when using any ordering studied, other than a magnitude ordering. This difference is almost a factor of 2 in many larger examples when using a lexicographic ordering. This encourages

Table 3. Equivalence of CNOT-Replaced Circuits

original circuit, central Z			
parity		central gate sequence	
even		R_z	
odd		$X R_z X$	
modified circuit, central Z			
exterior parity	inferior parity	overall parity	central gate sequence
even	even	even	R_z
even	odd	odd	$X R_z X$
odd	even	odd	$X R_z X$
odd	odd	even	$XX R_z XX = R_z$
original circuit, central Y			
parity		central gate sequence	
even		$Y R_z Y^\dagger$	
odd		$Y X R_z X Y^\dagger$	
modified circuit, central Z			
exterior parity	interior parity	overall parity	central gate sequence
even	even	even	$Y R_z Y^\dagger$
even	odd	odd	$Y X R_z X Y^\dagger$
odd	even	odd	$XY R_z Y^\dagger X = Y X R_z X Y^\dagger$
odd	odd	even	$X Y X R_z X Y^\dagger X = Y R_z Y^\dagger$

the use of the Bravyi–Kitaev mapping, as it could outweigh the relatively small benefit from the use of the Jordan–Wigner mapping in optimized lexicographic circuits.

Our results demonstrate that the performance of the Bravyi–Kitaev mapping is dependent on a variety of factors. While it is superior to the Jordan–Wigner mapping in most cases studied, several exceptions were observed. This emphasizes the importance of numerical analysis in future work. It is apparent that such studies must be performed across a range of molecular systems, with due consideration given to the region where classical full configuration interaction calculations are intractable.

Recent hardware developments have prompted the suggestion that quantum devices could be used for practical tasks in as little as five years.¹⁶ The use of quantum computers to perform classically intractable quantum chemistry calculations is often cited as one of the principal uses of emerging quantum technology.⁹

We have demonstrated here that the use of the Bravyi–Kitaev transformation frequently results in substantially reduced gate count estimates. In the future, we anticipate that the application of this mapping will assist in the performance of electronic structure calculations on real quantum computers, yielding results that have proven computationally elusive for classical devices.

APPENDICES

Details of the systems studied can be found in Table 2. Table 3 indicates the operation which is performed upon a qubit having modified the original circuit in line with section 6. Note that the end result of each modified circuit is equivalent to that of the original circuit.

AUTHOR INFORMATION

Corresponding Author

*E-mail: p.v.coveney@ucl.ac.uk

ORCID

Peter V. Coveney: 0000-0002-8787-7256

Funding

A.T. acknowledges EPSRC for a Ph.D. Studentship through Imperial College London's Centre for Doctoral Training in Controlled Quantum Dynamics (grant number EP/G037043/1). P.V.C. thanks EPSRC for funding via grant number EP/L00030X/1. P.J.L. acknowledges support by AFOSR award no. FA9550-12-1-0046.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The authors thank Jarrod McClean and Ryan Babbush for providing molecular orbital integral data used for preliminary versions of this analysis.

REFERENCES

- (1) Szabo, A.; Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*; Dover Publications: Mineola, NY, 1996.
- (2) Szabo, A.; Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*; Dover Publications: Mineola, NY, 1996; pp 231–270.
- (3) Gan, Z.; Grant, D. J.; Harrison, R. J.; Dixon, D. A. The Lowest Energy States of the Group-III–Group-VA Heteronuclear Diatomics: BN, BP, AlN, and AlP from Full Configuration Interaction Calculations. *J. Chem. Phys.* **2006**, *125*, 124311.
- (4) Aspuru-Guzik, A.; Dutoi, A. D.; Love, P. J.; Head-Gordon, M. Simulated Quantum Computation of Molecular Energies. *Science* **2005**, *309*, 1704–1707.
- (5) Lidar, D. A.; Wang, H. Calculating the Thermal Rate Constant with Exponential Speedup on a Quantum Computer. *Phys. Rev. E: Stat. Phys., Plasmas, Fluids, Relat. Interdiscip. Top.* **1999**, *59*, 2429–2438.
- (6) Kassal, I.; Aspuru-Guzik, A. Quantum Algorithm for Molecular Properties and Geometry Optimization. *J. Chem. Phys.* **2009**, *131*, 224102.
- (7) Reiher, M.; Wiebe, N.; Svore, K. M.; Wecker, D.; Troyer, M. Elucidating Reaction Mechanisms on Quantum Computers. *Proc. Natl. Acad. Sci. U. S. A.* **2017**, *114*, 7555–7560.
- (8) Kassal, I.; Jordan, S. P.; Love, P. J.; Mohseni, M.; Aspuru-Guzik, A. Polynomial-Time Quantum Algorithm for the Simulation of Chemical Dynamics. *Proc. Natl. Acad. Sci. U. S. A.* **2008**, *105*, 18681–18686.
- (9) Olson, J.; Cao, Y.; Romero, J.; Johnson, P.; Dallaire-Demers, P.-L.; Sawaya, N.; Narang, P.; Kivlichan, I.; Wasielewski, M.; Aspuru-Guzik, A. Quantum Information and Computation for Chemistry. 2017, *arXiv:1706.05413*. arXiv.org e-Print archive. <https://arxiv.org/abs/1706.05413>.
- (10) Lanyon, B. P.; Whitfield, J. D.; Gillett, G. G.; Goggin, M. E.; Almeida, M. P.; Kassal, I.; Biamonte, J. D.; Mohseni, M.; Powell, B. J.; Barbieri, M.; Aspuru-Guzik, A.; White, A. G. Towards Quantum Chemistry on a Quantum Computer. *Nat. Chem.* **2010**, *2*, 106–111.
- (11) Du, J.; Xu, N.; Peng, X.; Wang, P.; Wu, S.; Lu, D. NMR Implementation of a Molecular Hydrogen Quantum Simulation with Adiabatic State Preparation. *Phys. Rev. Lett.* **2010**, *104*, 030502.
- (12) O'Malley, P. J. J.; Babbush, R.; Kivlichan, I. D.; Romero, J.; McClean, J. R.; Barends, R.; Kelly, J.; Roushan, P.; Tranter, A.; Ding, N.; Campbell, B.; Chen, Y.; Chen, Z.; Chiaro, B.; Dunsworth, A.; Fowler, A. G.; Jeffrey, E.; Lucero, E.; Megrant, A.; Mutus, J. Y.; Neeley, M.; Neill, C.; Quintana, C.; Sank, D.; Vainsencher, A.; Wenner, J.; White, T. C.; Coveney, P. V.; Love, P. J.; Neven, H.; Aspuru-Guzik, A.; Martinis, J. M. Scalable Quantum Simulation of Molecular Energies. *Phys. Rev. X* **2016**, *6*, 031007.
- (13) McClean, J. R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. The Theory of Variational Hybrid Quantum-Classical Algorithms. *New J. Phys.* **2016**, *18*, 023023.

- (14) Kandala, A.; Mezzacapo, A.; Temme, K.; Takita, M.; Brink, M.; Chow, J. M.; Gambetta, J. M. Hardware-Efficient Variational Quantum Eigensolver for Small Molecules and Quantum Magnets. *Nature* **2017**, *549*, 242–246.
- (15) Sete, E. A.; Zeng, W. J.; Rigetti, C. T. Functional Architecture for Scalable Quantum Computing. *2016 IEEE International Conference on Rebooting Computing (ICRC)* **2016**, 1–6.
- (16) Mohseni, M.; Read, P.; Neven, H.; Boixo, S.; Denchev, V.; Babbush, R.; Fowler, A.; Smelyanskiy, V.; Martinis, J. Commercialize Quantum Technologies in Five Years. *Nature* **2017**, *543*, 171.
- (17) Jordan, P.; Wigner, E. Über das Paulische Äquivalenzverbot. *Eur. Phys. J. A* **1928**, *47*, 631–651.
- (18) Bravyi, S. B.; Kitaev, A. Y. Fermionic Quantum Computation. *Ann. Phys.* **2002**, *298*, 210–226.
- (19) Seeley, J. T.; Richard, M. J.; Love, P. J. The Bravyi-Kitaev Transformation for Quantum Computation of Electronic Structure. *J. Chem. Phys.* **2012**, *137*, 224109.
- (20) Tranter, A.; Sofia, S.; Seeley, J.; Kaicher, M.; McClean, J.; Babbush, R.; Coveney, P. V.; Mintert, F.; Wilhelm, F.; Love, P. J. The Bravyi-Kitaev Transformation: Properties and Applications. *Int. J. Quantum Chem.* **2015**, *115*, 1431–1441.
- (21) Trotter, H. F. On the Product of Semi-Groups of Operators. *Proc. Am. Math. Soc.* **1959**, *10*, 545–551.
- (22) Suzuki, M. Generalized Trotter's Formula and Systematic Approximants of Exponential Operators and Inner Derivations with Applications to Many-Body Problems. *Commun. Math. Phys.* **1976**, *51*, 183–190.
- (23) Hatano, N.; Suzuki, M. In *Quantum Annealing and Other Optimization Methods*; Das, A., Chakrabarti, B. K., Eds.; Springer: Berlin, Heidelberg, 2005; pp 37–68.
- (24) Kitaev, A. Y. Quantum Measurements and the Abelian Stabilizer Problem. 1995, *arXiv:quant-ph/9511026*. arXiv e-Print archive. <https://arxiv.org/abs/quant-ph/9511026>.
- (25) McClean, J. R.; Babbush, R.; Love, P. J.; Aspuru-Guzik, A. Exploiting Locality in Quantum Computation for Quantum Chemistry. *J. Phys. Chem. Lett.* **2014**, *5*, 4368–4380.
- (26) Havlíček, V.; Troyer, M.; Whitfield, J. D. Operator Locality in Quantum Simulation of Fermionic Models. *Phys. Rev. A: At, Mol, Opt. Phys.* **2017**, *95*, 032332 DOI: 10.1103/PhysRevA.95.032332.
- (27) Setia, K.; Whitfield, J. D. Bravyi-Kitaev Superfast Simulation of Fermions on a Quantum Computer. 2017, *arXiv:1712.00446*. arXiv e-Print archive. <https://arxiv.org/abs/1712.00446>.
- (28) Helgaker, T.; Jørgensen, P.; Olsen, J. *Molecular Electronic-Structure Theory*; Wiley: Chichester, New York, 2000.
- (29) Somma, R.; Ortiz, G.; Knill, E.; Gubernatis, J. Quantum Simulations of Physics Problems. *Int. J. Quantum Inf.* **2003**, *01*, 189–206.
- (30) This requirement can be relaxed by additional encoding techniques;^{46,47} however, we do not consider these techniques here.
- (31) Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; O'Brien, J. L. A Variational Eigenvalue Solver on a Photonic Quantum Processor. *Nat. Commun.* **2014**, *5*, 4213.
- (32) Poulin, D.; Hastings, M. B.; Wecker, D.; Wiebe, N.; Doherty, A. C.; Troyer, M. The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry. *Quantum Inf Comput* **2015**, *15*, 361–384.
- (33) Babbush, R.; McClean, J.; Wecker, D.; Aspuru-Guzik, A.; Wiebe, N. Chemical Basis of Trotter-Suzuki Errors in Quantum Chemistry Simulation. *Phys. Rev. A: At, Mol, Opt. Phys.* **2015**, *91*, 022311.
- (34) Whitfield, J. D.; Biamonte, J.; Aspuru-Guzik, A. Simulation of Electronic Structure Hamiltonians Using Quantum Computers. *Mol. Phys.* **2011**, *109*, 735–750.
- (35) Johnson, R. D., III NIST Computational Chemistry Comparison and Benchmark Database NIST Standard Reference Database Number 101 Release 18; 2016 See the following: <https://cccbdb.nist.gov/>.
- (36) Parrish, R. M.; Burns, L. A.; Smith, D. G. A.; Simmonett, A. C.; DePrince, A. E.; Hohenstein, E. G.; Bozkaya, U.; Sokolov, A. Y.; Di Remigio, R.; Richard, R. M.; Gonthier, J. F.; James, A. M.; McAlexander, H. R.; Kumar, A.; Saitow, M.; Wang, X.; Pritchard, B. P.; Verma, P.; Schaefer, H. F.; Patkowski, K.; King, R. A.; Valeev, E. F.; Evangelista, F. A.; Turney, J. M.; Crawford, T. D.; Sherrill, C. D. Psi4 1.1: An Open-Source Electronic Structure Program Emphasizing Automation, Advanced Libraries, and Interoperability. *J. Chem. Theory Comput.* **2017**, *13*, 3185–3197.
- (37) McClean, J. R.; Kivlichan, I. D.; Steiger, D. S.; Sung, K. J.; Cao, Y.; Dai, C.; Fried, E. S.; Gidney, C.; Häner, T.; Havlíček, V.; Huang, C.; Jiang, Z.; Neeley, M.; Romero, J.; Rubin, N.; Sawaya, N. P. D.; Setia, K.; Sim, S.; Sun, W.; Zhang, F.; Babbush, R. OpenFermion: The Electronic Structure Package for Quantum Computers. 2017, *arXiv:1710.07629*. arXiv e-Print archive. <https://arxiv.org/abs/1710.07629>.
- (38) Hastings, M. B.; Wecker, D.; Bauer, B.; Troyer, M. Improving Quantum Algorithms for Quantum Chemistry. *Quantum Inf. Comput.* **2015**, *15*, 1–21.
- (39) Häner, T.; Steiger, D. S. 0.5 Petabyte Simulation of a 45-Qubit Quantum Circuit. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* **2017**, 33:1–33:10.
- (40) Pednault, E.; Gunnels, J. A.; Nannicini, G.; Horesh, L.; Magerlein, T.; Solomonik, E.; Wisnieff, R. Breaking the 49-Qubit Barrier in the Simulation of Quantum Circuits. 2017, *arXiv:1710.05867*. arXiv e-Print archive. <https://arxiv.org/abs/1710.05867>.
- (41) Smelyanskiy, M.; Sawaya, N. P. D.; Aspuru-Guzik, A. qHiPSTER: The Quantum High Performance Software Testing Environment. 2016, *arXiv:1601.07195* arXiv e-Print archive. <https://arxiv.org/abs/1601.07195>.
- (42) Nielsen, M. A., Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*; Cambridge University Press, 2010.
- (43) Rossi, E.; Bendazzoli, G. L.; Evangelisti, S.; Maynau, D. A Full-Configuration Benchmark for the N2Molecule. *Chem. Phys. Lett.* **1999**, *310*, 530–536.
- (44) This iterative process is not optimized for classical computing cost, and would likely be performed more efficiently when generating circuits to be implemented on real quantum devices.
- (45) Jones, E.; Oliphant, T.; Peterson, P. et al. *SciPy: Open Source Scientific Tools for Python*. 2001, Online (accessed May 2, 2018).
- (46) Bravyi, S.; Gambetta, J. M.; Mezzacapo, A.; Temme, K. Tapering off Qubits to Simulate Fermionic Hamiltonians. 2017, *arXiv:1701.08213*. arXiv e-Print archive. <https://arxiv.org/abs/1701.08213>.
- (47) Steudtner, M.; Wehner, S. Lowering Qubit Requirements for Quantum Simulations of Fermionic Systems. 2017, *arXiv:1712.07067*. arXiv e-Print archive. <https://arxiv.org/abs/1712.07067>.