Ordering of Trotterization: Impact on Errors in Quantum Simulation of Electronic Structure

Andrew Tranter [1,*], Peter J. Love [1], Florian Mintert [2], Nathan Wiebe [3,4] and Peter V. Coveney [5]

[1]   Department of Physics and Astronomy, Tufts University, Medford, MA 02155, USA; peter.love@tufts.edu
[2]   Department of Physics, Imperial College London, London SW7 2AZ, UK; f.mintert@imperial.ac.uk
[3]   Department of Physics, University of Washington, Seattle, WA 98105, USA; nwiebe@uw.edu
[4]   Pacific Northwest National Laboratory, Richland, WA 98382, USA
[5]   Centre for Computational Science, University College London, London WC1H 0AJ, UK;
       p.v.coveney@ucl.ac.uk
*   Correspondence: andrew.tranter@tufts.edu

Abstract: Trotter–Suzuki decompositions are frequently used in the quantum simulation of quantum chemistry. They transform the evolution operator into a form implementable on a quantum device, while incurring an error—the Trotter error. The Trotter error can be made arbitrarily small by increasing the Trotter number. However, this increases the length of the quantum circuits required, which may be impractical. It is therefore desirable to find methods of reducing the Trotter error through alternate means. The Trotter error is dependent on the order in which individual term unitaries are applied. Due to the factorial growth in the number of possible orderings with respect to the number of terms, finding an optimal strategy for ordering Trotter sequences is difficult. In this paper, we propose three ordering strategies, and assess their impact on the Trotter error incurred. Initially, we exhaustively examine the possible orderings for molecular hydrogen in a STO-3G basis. We demonstrate how the optimal ordering scheme depends on the compatibility graph of the Hamiltonian, and show how it varies with increasing bond length. We then use 44 molecular Hamiltonians to evaluate two strategies based on coloring their incompatibility graphs, while considering the properties of the obtained colorings. We find that the Trotter error for most systems involving heavy atoms, using a reference magnitude ordering, is less than 1 kcal/mol. Relative to this, the difference between ordering schemes can be substantial, being approximately on the order of millihartrees. The coloring-based ordering schemes are reasonably promising—particularly for systems involving heavy atoms—however further work is required to increase dependence on the magnitude of terms. Finally, we consider ordering strategies based on the norm of the Trotter error operator, including an iterative method for generating the new error operator terms added upon insertion of a term into an ordered Hamiltonian.

Keywords: quantum simulation; electronic structure theory; quantum computing

1. Introduction

Computational chemistry is the use of computers to answer outstanding questions in chemistry. Various algorithms have been developed to calculate the properties of molecules and reactions. The predictions of these methods are used in many fields. For example, they are frequently used to aid development of synthetic processes [1] and target searches in drug discovery [2]. Similarly, they can be used to characterize molecular configurations which are difficult to study experimentally, such as transition states of chemical reactions.

The initial step of most computational approaches to chemistry involves some form of electronic structure theory calculation: the determination of molecular electronic wavefunctions and their corresponding energies [3,4]. Many methods for the solution of this problem have been established.

The most conceptually simple approach within this category is the full configuration interaction (FCI) approach. Here, a finite basis set of spin-orbitals is used to describe the Hilbert space of the electronic wavefunction. Typically, this is initially a localized basis of atomic orbitals. Molecular orbitals are found using an algorithm such as the Hartree–Fock method, with the Hamiltonian in a basis of Slater determinants formed from these molecular orbitals giving the configuration interaction (CI) matrix. The eigenstates and eigenvalues of the CI matrix then give electronic eigenstates and their corresponding energies. Such a technique is numerically exact, to within the limitations of the basis set and the assumption that relativistic effects are negligible. However, the computational expense of this technique scales factorially with the number of basis functions used [3]. This limits the application to extremely small molecules [5], and thus is typically used as a benchmark for other methods.

Approximate methods such as coupled cluster theory [6] or Møller–Plesset perturbation theory [7] are often used to obtain results with practical computational resources; however, exact methods are required for benchmarking these. Additionally, for any approximation applied, a system can be found wherein such an approximation breaks down. These facts reinforce the necessity for computationally feasible numerically exact methods.

It has been established that a scalable quantum computer would be capable of providing full configuration interaction level electronic structure results in polynomial time [8–10]. As quantum chemistry is expected to be a key application of developing quantum devices [11], there has been a great deal of theoretical development on the algorithms that would be used to perform quantum chemical calculations on a scalable quantum computer. Algorithms to describe various chemical processes have been developed, including energy spectra [8], reaction rates [12–14] and reaction dynamics [15]. Experimental demonstrations have been shown on a variety of quantum computing architectures, including photonic [16], nuclear magnetic resonance [17], superconducting [18,19] and trapped-ion [20,21] systems.

The variational quantum eigensolver (VQE) algorithm [22,23] is a hybrid quantum-classical scheme, where a classical variational approach is taken with a quantum device used to determine accurate expectation values. This approach has obtained much attention in the recent literature, although other hybrid quantum-classical schemes for the simulation of quantum systems [24,25] and other purposes [26–28] have been reported. Although repeated Ansatz preparation and the need for multiple variational steps increases the overall asymptotic cost of these algorithms, they do not require coherency to be maintained throughout the entire circuit, instead requiring many short coherent evolutions. As such, while not providing full configuration interaction level accuracy, these methods can efficiently provide results that are more accurate than classical equivalents. These methods allowed for the simulation of beryllium hydride in a minimal basis in 2017, albeit with low accuracy [19].

Although recent experimental approaches have focused on the variational quantum eigensolver, open questions remain in the study of approaches based on a phase estimation algorithm [8,29–31]. This requires coherency to be maintained for a time that scales exponentially with the desired precision. While this is not asymptotically prohibitive due to the required precision being fixed, it is markedly less practical for implementation on noisy devices. However, this does result in eigenvalues which are numerically exact.

Trotterization is the use of Trotter–Suzuki formulae to simulate evolution under a given Hamiltonian which comprises a sum of many terms, by sequentially simulating the evolution under each term. In both phase estimation and VQE, Trotterization of the qubit Hamiltonian is a basic tool. It can be used for the time evolution required by phase estimation and for Ansatz preparation in VQE. In doing this, a degree of error—the *Trotter error*—is introduced. This error can be made arbitrarily small by increasing the Trotter number, i.e., the number of times the terms are iterated through. However, this increases the quantum computational cost of the procedure. As it is likely that

quantum computational resources will be highly limited in the foreseeable future, it is useful to find methods to reduce this overhead in order to enable the simulation of larger systems. This is particularly relevant when considering phase estimation approaches, as Trotter error in these translates to the only source of algorithmic error in the simulation, whereas the impact of this error in a variational scheme is minimized through optimization of Ansatz parameters [32].

It has been shown [33,34] that the error incurred in the use of a Trotter–Suzuki approximation is dependent on the order in which individual terms are simulated. However, the use of a Trotter ordering scheme is not determined by the Trotter error effects alone. Differing Trotter ordering can have an effect on the length of the circuit required to simulate one Trotter step, as identical gates between terms can be canceled [33], with the degree of cancellation being dependent on the ordering of terms. Several ordering schemes have been reported to this end [33–36]. Nonetheless, it is possible that this effect could be outweighed by the minimization of Trotter error, and the consequent reduction in the overall number of Trotter steps—especially where the Trotter number necessary for chemically accurate predictions is low, or the ordering impact on Trotter error is high. In this paper, we thus focus on attempts to reduce the single-Trotter-step Trotter error, instead of considering other effects that impact circuit length. To this end, we report and characterize two new ordering strategies.

It is possible to derive an analytic expression for the *Trotter error operator* [35], of which the expectation value in the ground state yields the Trotter error. As finding the exact ground state is exponentially hard, finding the exact Trotter error with this approach is similarly difficult. However, the norm of the Trotter error operator can be determined without diagonalizing the full Hamiltonian. This serves as an upper bound to the true Trotter error, although this bound can be extremely loose [36]. It has been observed [36] that, despite the looseness of this bound, it does replicate some qualitative trends of the true Trotter error. As such, it is possible that using this information to guide an ordering scheme could result in an effective strategy. In Section 5, we report and assess an ordering scheme in this vein. While the Trotter error operator can be computed in polynomial time, it remains computationally intensive. We therefore derive a Trotter term insertion error operator, which yields the terms that are added to the Trotter error operator upon insertion of a new term into the Hamiltonian. This allows direct and efficient comparison of the impact upon the Trotter error operator, of inserting a term into the Hamiltonian in different positions. This is of use in the greedy ordering scheme discussed in Section 5.

We begin by providing a brief overview of the theory underpinning the canonical methods of the quantum simulation of electronic structure theory, and particularly the theory of Trotterization [8]. We then consider three approaches to the development of Trotter ordering schemes. Firstly, in Section 3, we discuss the hydrogen molecule in a minimal basis as a simple case study. Although this system has been extensively studied, including the consideration of all possible Trotter orderings in an experimental context [18], an examination of the distribution of Trotter errors across varying Trotter orderings is yet unreported in the literature. We consider this here, along with a discussion of the geometry dependence of the optimal Trotter ordering. Secondly, in Section 4, we propose two ordering schemes based on subdividing the molecular Hamiltonian into mutually commuting subsets of terms, and report their performance across a dataset of 44 systems. Finally, in Section 5, we propose and assess a final ordering scheme based on minimizing the norm of the Trotter error operator.

2. Trotterization—Theoretical Background

The electronic Hamiltonian in the second quantized formalism is given by:

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

where h_{ij} and h_{ijkl} are Coulombic overlap and exchange integrals determined by the basis set chosen [3,4]. The overall goal of the simulation process is to determine the eigenstates and corresponding eigenvalues of this Hamiltonian. Letting the number of spin-orbitals included in

the basis set be N, there are $O(N^4)$ terms in the Hamiltonian in Equation (1). In practice, molecular symmetries and orbital localization result in many of these terms having a zero or negligible coefficient. Regardless, the determination of the h_{ij} and h_{iijkl} integrals is classically efficient, and thus these weighting constants can be considered as input data when performing the calculation on either a classical or quantum device.

Classically, the matrix elements of the electronic Hamiltonian in a basis of Slater determinants may be obtained. The eigenvalues of the full configuration interaction matrix can then be determined, for instance by direct diagonalization. However, the dimension of the Fock space which $\hat{H}$ acts upon grows exponentially with N. This can be reduced by excluding the subspace with the incorrect number of electrons; however, this still results in a growth of $\binom{N}{n}$, where n is the number of electrons. As such, it is intractable to perform this process for more than a handful of tens of spin-orbitals. Conversely, N qubits span the entire Fock space that $\hat{H}$ acts upon—a quantum computer thus circumvents the exponential cost of the simulation procedure.

Performing a similar procedure on a quantum device differs from a classical full configuration interaction calculation. Having calculated the h_{ij} and h_{ijkl} integrals, a mapping scheme to transform the creation and annihilation operators of Equation (1), along with the electronic states they act upon, to operations upon and states of qubits must be found. This is typically performed through the Jordan–Wigner transformation, the Bravyi–Kitaev transformation, or other constructions [34,37–41]. In this paper, we primarily utilize the Jordan–Wigner transformation, although we present some results using the Bravyi–Kitaev mapping in Section 4. The result of performing this mapping procedure is the generation of a qubit Hamiltonian, which consists of a sum of weighted strings of qubit Pauli operators.

From here, two approaches to finding molecular eigenstates and eigenvalues are common. In quantum phase estimation [29,42], a circuit corresponding to the the evolution operator $U = \exp\left(-it\hat{H}\hbar\right)$ of the qubit Hamiltonian is repeatedly applied, with a second register of qubits used to store a binary expansion of the true eigenvalue. In a variational quantum eigensolver, a parameterized Ansatz state which is classically hard to store is generated. The expectation value of the qubit Hamiltonian with this state is measured, and a classical optimizer used to vary the Ansatz parameters until the expectation value is variationally minimized [22,23]. Typically, the unitary coupled cluster Ansatz [32]—or a form derived from it [43]—is used. Here, an exponentiated form of the unitary cluster operator is applied to a reference state.

In both cases, a circuit that implements an exponentiated sum of Pauli operators must be found. In general, such a circuit is difficult to find. However, standard circuits exist to simulate individual Pauli strings—$U_i = \exp\left(-it_i\hat{H}_i\hbar\right)$. As such, we invoke a Trotter–Suzuki approximation to break the overall evolution operator into a product of individual terms.

A first-order Trotter–Suzuki approximation is given by [44]:

$$e^{-i\hat{H}t} \approx \left(\prod_{k=1}^{m} e^{-i\hat{H}_k t/N_T}\right)^{N_T}. \tag{2}$$

where the Hamiltonian $\hat{H}$ is a sum of m terms $\hat{H}_k$. Intuitively, this states that the Hamiltonian can be approximated by rapidly switching between each term, across the desired evolution time. The exponential of a single Hamiltonian term can be described easily. Increasing N_T, the *Trotter number*—the number of Trotter steps—deterministically reduces the error incurred by this approximation, but with commensurate increase in circuit length.

A second-order Trotter–Suzuki approximant

$$e^{-i\hat{H}t} \approx \left(\prod_{k=1}^{m} e^{-i\hat{H}_k t/2N_T} \prod_{k=m}^{1} e^{-i\hat{H}_k t/2N_T}\right)^{N_T} \tag{3}$$

can be obtained by symmetrizing the first-order approximant. Higher-order approximants can be obtained recursively [45]. Increasing the approximation order will yield reductions in error. However,

the overall length of the quantum circuit required to simulate higher-order Trotter approximants increases exponentially with the order used. Choosing an appropriate strategy for simulation is then a three-way trade-off among the order of Trotter–Suzuki approximant, the number of time steps, and the need to minimize the overall circuit length. For consistency with other work, we mostly restrict ourselves here to use of a second-order Trotter approximant (Equation (3)), although we additionally use a first-order Trotter approximant in Section 3.

　　To bound the error incurred in the use of this approximation, Poulin et al. [35] introduced the Trotter *error operator*, which gives the Trotter error for a given Trotterized Hamiltonian. For a second-order Trotter–Suzuki approximant, the expectation value of this with an eigenstate is given by [36]:

$$\Delta E_i = -\frac{\Delta_t^2}{12} \sum_{\alpha \leq \beta} \sum_{\beta} \sum_{\gamma < \beta} \langle \psi_i | \left[\hat{H}_\alpha \left(1 - \frac{\delta_{\alpha,\beta}}{2} \right), [\hat{H}_\beta, \hat{H}_\gamma] \right] | \psi_i \rangle, \tag{4}$$

where ΔE_i is the expected error (ignoring higher-order terms), Δ_t^2 is the Trotter step size, $|\psi_i\rangle$ is the ground state , δ is the Kronecker delta function and $\hat{H}_i$ are Hamiltonian terms [36].

　　Examining this, it is evident that two factors primarily affect the degree of error introduced by the use of the Trotter–Suzuki approximation. Firstly, the error will decrease with the number of time steps used. Additionally, as the sums in Equation (4) are ordered—being sums over conditional indices, rather than the whole range—the order in which terms are applied will impact the error incurred. Indeed, previous work has suggested that the Trotter ordering can dramatically impact the number of Trotter steps required for constant precision [33,39]. It is thus desirable to find an ordering strategy which minimizes this error, motivating this paper.

　　In this section, we have discussed the underlying theory of Trotterization and the impact of Trotter ordering schemes. For most molecular systems, the space of possible orderings is sufficiently vast as to inhibit direct statistical analysis. We proceed to consider a small test case—molecular hydrogen in a minimal basis—both to justify the generalized ordering schemes presented in Section 4 and to study the role of molecular geometry in a small example.

3. Molecular Hydrogen

　　Molecular hydrogen in an STO-3G basis is the smallest and simplest chemically interesting electronic structure problem—barring its cationic form, which lacks two-electron interactions. Due to its simplicity, this example has been widely used both experimentally and theoretically [8,16–18,20,39,41]. It was recently shown that this example is in a sense not quantum mechanical, as it lacks measurement contextuality [46]; however, the simplicity of the system makes it a good candidate for initial studies.

　　The qubit Hamiltonian for the hydrogen molecule in an STO-3G basis, using a Jordan–Wigner transformation, with a bond length of 0.7414 Å, is given by:

$$
\begin{aligned}
\hat{H}_q = &-0.81262I + 0.17120\sigma_0^z + 0.17120\sigma_1^z - 0.22279\sigma_2^z \\
&-0.22279\sigma_3^z + 0.16862\sigma_1^z\sigma_0^z + 0.12054\sigma_2^z\sigma_0^z + 0.16587\sigma_3^z\sigma_0^z \\
&+0.16587\sigma_2^z\sigma_1^z + 0.12054\sigma_3^z\sigma_1^z + 0.17435\sigma_3^z\sigma_2^z \\
&-0.04532\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^x + 0.04532\sigma_3^x\sigma_2^y\sigma_1^y\sigma_0^x \\
&+0.04532\sigma_3^y\sigma_2^x\sigma_1^x\sigma_0^y - 0.04532\sigma_3^x\sigma_2^x\sigma_1^y\sigma_0^y .
\end{aligned}
\tag{5}
$$

where the lower index of each Pauli operator corresponds to the index of the qubit it is applied to. This consists of 15 terms. Unfortunately, even a Hamiltonian of this small size provides on the order of 10^{12} possible orderings. It is thus infeasible to consider all orderings for the full Hamiltonian, even for the hydrogen molecule in a minimal basis. The second-order Trotter error operator (Equation (4)) consists of a sum of many triple commutators between Hamiltonian terms. As such, it is prudent

to consider the commutativity of the terms in the Hamiltonian. Inspecting Equation (5) shows that the terms with only two Pauli Z operators commute with all terms in the Hamiltonian. Along with the identity term, these form a *totally commuting set*. Because they commute with all terms, their exponentials commute with the exponentials of all other terms, and consequentially they can be freely moved around a Trotterized evolution operator. As such, all terms in this set can be moved to the front of the operator, combined, and simulated as an entirely independent operator

$$ e^{\frac{-i\hat{H}t}{\hbar}} \approx e^{\frac{-i\hat{H}_C}{\hbar}} T\left(\hat{H}_A\right) , \tag{6} $$

where $\hat{H}_C$ is the sum of all totally commuting terms, and $T\left(\hat{H}_A\right)$ is a Trotterization of the other terms. In other words, we only need to Trotterize the terms that do not commute with all terms.

This analysis has two principal advantages. Firstly, as the totally commuting set does not require a Trotter–Suzuki approximation, these terms can be simulated in one time step. This dramatically cuts down on the number of gates required for implementation on a quantum device. It should be noted here that for larger systems the size of the totally commuting set rapidly drops to one term (the identity operator). As such, this advantage is not scalable, and therefore this technique cannot be used directly in the analysis of larger systems.

More importantly for the purposes of this discussion, as we do not need to Trotterize the terms in the totally commuting set, the space of possible orderings is dramatically reduced to $40,320$. As the simulation of each ordering takes less than 1 s on an average laptop computer, it is thus feasible to simply brute force search the entire space of orderings in order to study the distribution of errors.

In our simulations, integral data were generated using the Psi4 [47] quantum chemistry package and OpenFermion [48], with a standard Hartree–Fock basis used to express the molecular Hamiltonian. Our Python code was used to generate both Jordan–Wigner and Bravyi–Kitaev Hamiltonians. The exact ground state and its corresponding energy were then determined. Similarly, the Trotter error for every possible ordering was determined, with the overall evolution time and the number of Trotter steps set to unity.

Figure 1 is a cumulative density plot showing the distribution of Trotterization errors which are obtained as a result of ordering variations. Errors are given relative to the true eigenvalue of the Hamiltonian with totally commuting terms removed. The same distribution is obtained regardless of whether the Jordan–Wigner or Bravyi–Kitaev Hamiltonian is simulated. The distribution is heavily weighted towards the low error region. This is promising, as it implies that a random choice of ordering is likely to introduce a small amount of ordering-dependent error. There is, however, over an order of magnitude difference between the optimal and poorest orderings. In this case, the difference is such that only three Trotter steps are required to reach an accuracy of 0.0001 a.u. for the optimal ordering, whereas seven Trotter steps are required for the worst ordering. This implies that—if we extrapolate solely from the hydrogen molecule in a minimal basis—picking a bad ordering could result in a dramatically increased gate count in actual quantum simulations. This is important despite the high likelihood of random orderings being accurate, as if a systematic approach to ordering is taken, it must be ensured that this does not result in bad orderings being selected. It should be noted here, however, that the results of Section 4 show that this difference is lessened in systems involving heavier atoms.
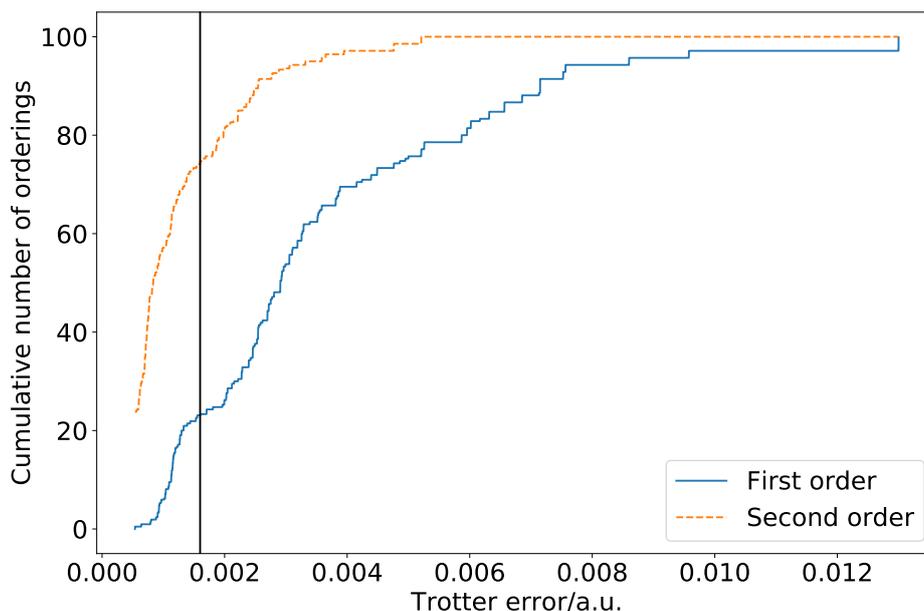
Figure 1. Cumulative density plot of hydrogen ordering errors for one Trotter step, using a bond length of 0.7414 Å and an STO-3G atomic basis. The vertical line denotes an error of 1 kcal/mol. Approximately 20% of the orderings achieve this error or lower for the first-order Trotter–Suzuki approximation. Around 80% of orderings achieve an error of 0.005 Hartree, approximately half that of the worst possible ordering.

We now consider the optimal ordering strategy for the hydrogen molecule, by re-examining the commutativity structure graph (the incompatibility graph) discussed above. Figure 2 shows this with totally commuting terms removed. Here, it is apparent that the incompatibility graph is bipartite. As such, the Hamiltonian can be subdivided into two sets of terms. Within each set, all terms mutually commute. No term in a given set commutes with any term in the other set. One set consists of all of the terms with Pauli Z operators, whereas the other set consists of terms with Pauli X and Y operators. We label the former of these sets the Z-set and the latter the XY-set. Using the Jordan–Wigner transformation, the optimal ordering using a first-order Trotter expansion is given by

$$\begin{aligned}
\hat{H}_1 =\; &0.04532\sigma_3^x\sigma_2^y\sigma_1^y\sigma_0^x - 0.22280\sigma_2^z - 0.04532\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^x \\
&- 0.22280\sigma_3^z - 0.04532\sigma_3^x\sigma_2^x\sigma_1^y\sigma_0^y + 0.17120\sigma_1^z \\
&+ 0.04532\sigma_3^y\sigma_2^x\sigma_1^x\sigma_0^y + 0.17120\sigma_0^z \,,
\end{aligned}$$

where the leftmost term is simulated first, the second term is simulated second, and so on. Here, the first, third, fifth and seventh terms are members of the XY-set, while the remaining terms are members of the Z-set. Examining this, it is evident that the optimal ordering strategy is given by alternating between completely commuting sets in order of descending coefficient magnitude. This strategy is optimal for both Jordan–Wigner and Bravyi–Kitaev Hamiltonians. However, other systematic approaches to ordering strategies for this Hamiltonian produce differing results for the two mapping techniques. This emphasizes the need for a systematic ordering scheme which is mapping agnostic.
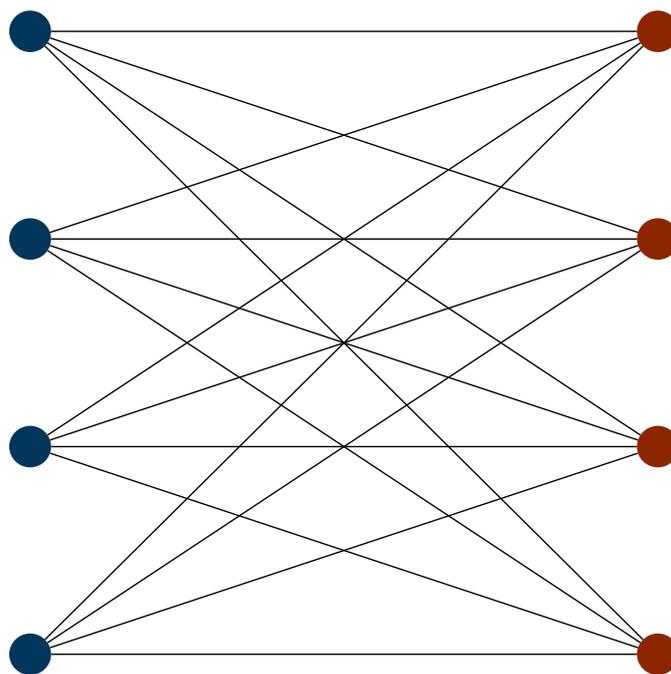
Figure 2. The incompatibility graph of the Jordan–Wigner and Bravyi–Kitaev Hamiltonians, with the totally commuting set removed. Nodes correspond to Hamiltonian terms, edges correspond to non-commutativity between terms. Two independent sets are clearly revealed, with the XY-set colored blue and the Z-set colored red.

Geometry Dependence

Molecular hydrogen allows for the specification of the entire molecular geometry with a single parameter; the bond length of the molecule. This allows for the consideration of how the Trotter error varies with the molecular geometry. Figure 3 demonstrates this. To contrast the results above and for consistency with prior work, we here use a second-order Trotter expansion. Increasing the bond length results in a substantially reduced Trotter error. This is likely due to the increased locality of the electronic eigenstates, resulting in increased degeneracy, and thus in increased symmetry in the coefficients of the Z-set terms. At asymptotic separation, the Z-set terms have equal coefficient. This additional symmetry reduces the dependence of the Trotter error on ordering choice. It also allows for increased cancellation in the terms of the Trotter error operator, resulting in a reduction in Trotter error overall at higher bond lengths.

Table 1 shows how the Pauli Hamiltonian and optimal Trotter ordering varies with bond length. We first observe that the optimal ordering has changed at equilibrium bond length, due to the use of the second-order Trotter–Suzuki formula. The strategy of alternating between commuting sets remains. However, within each set, the terms are ordered in descending magnitude order, likely due to the symmetrization of the second-order approximant.

For small bond length, we observe substantially increased difference in the coefficients of the Z-set terms, due to the increased relative stability of the bonding orbitals. The strategy of alternating between Z-set and XY-set terms is optimal at all pre-asymptotic separations. However, the particular choice of Z-set terms differs at higher bond lengths, due to the increasing similarity of the coefficients of the Z-set terms. At asymptotic separation, a different ordering is preferable. Here, as the Z-set terms are equal, the Trotter error can be reduced to zero by placing all XY-set terms at the start of the expansion.
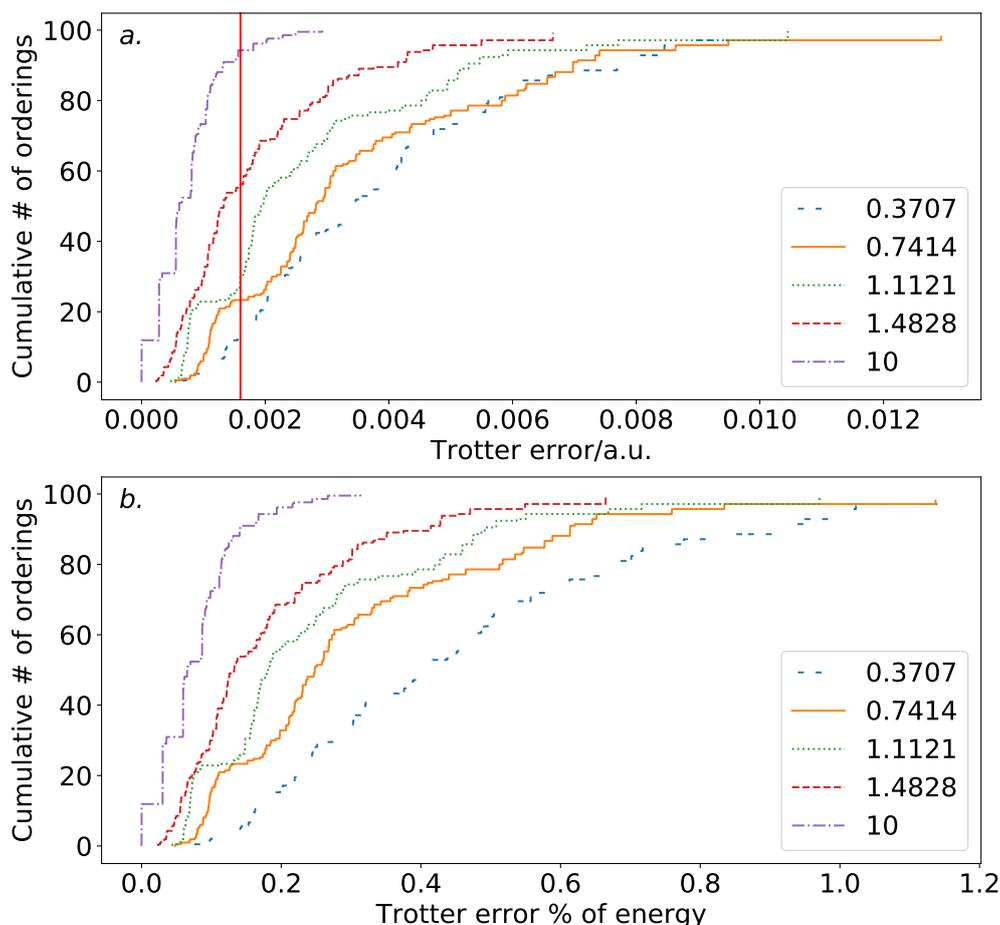
Figure 3. Distribution of Trotter errors by ordering for varying bond lengths for H_2 in a minimal basis. (**a**): versus absolute Trotter error. As the bond length decreases, both the Trotter error and the dependence of the Trotter error on the ordering chosen increase. (**b**): versus the Trotter error as a percentage of the ground state energy. The same trend as with the absolute Trotter error is observed, although the ordering dependence at extremely low bond length is accentuated.

Table 1. Optimal Trotter orderings for the Hydrogen molecule in a STO-3G basis, for varying bond length, using a second-order Trotter–Suzuki approximation and a Jordan–Wigner mapping. Each ordering proceeds from top to bottom. The ordering changes as bond length increases, although, prior to the asymptotic limit, all orderings are of the form of alternating between commuting sets. At asymptotic separation, it is preferable to simulate sets in sequence, due to the symmetry of the coefficients.

Bond length (Å)				
0.3707	**0.7414**	**1.1121**	**1.4828**	**10.000**
$0.24197\,\sigma_1^z$	$0.17120\,\sigma_1^z$	$-0.10205\,\sigma_2^z$	$-0.03780\,\sigma_2^z$	$0.09021\,\sigma_3^y\sigma_2^x\sigma_1^x\sigma_0^y$
$0.04084\,\sigma_3^y\sigma_2^x\sigma_1^x\sigma_0^y$	$0.04532\,\sigma_3^y\sigma_2^x\sigma_1^x\sigma_0^y$	$0.05100\,\sigma_3^y\sigma_2^x\sigma_1^x\sigma_0^y$	$0.05711\,\sigma_3^y\sigma_2^x\sigma_1^x\sigma_0^y$	$0.09021\,\sigma_3^x\sigma_2^y\sigma_1^y\sigma_0^x$
$0.24197\,\sigma_0^z$	$0.17120\,\sigma_0^z$	$-0.10205\,\sigma_3^z$	$-0.03780\,\sigma_3^z$	$-0.09021\,\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^y$
$-0.04084\,\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^x$	$-0.04532\,\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^x$	$0.05100\,\sigma_3^x\sigma_2^y\sigma_1^y\sigma_0^x$	$0.05711\,\sigma_3^x\sigma_2^y\sigma_1^y\sigma_0^x$	$-0.09021\,\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^x$
$-0.48079\,\sigma_3^z$	$-0.22279\,\sigma_3^z$	$0.12533\,\sigma_0^z$	$0.09462\,\sigma_0^z$	$0.03964\,\sigma_0^z$
$-0.04084\,\sigma_3^x\sigma_2^x\sigma_1^y\sigma_0^y$	$-0.04532\,\sigma_3^x\sigma_2^x\sigma_1^y\sigma_0^y$	$-0.05100\,\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^x$	$-0.05711\,\sigma_3^y\sigma_2^y\sigma_1^x\sigma_0^x$	$0.03964\,\sigma_1^z$
$-0.48079\,\sigma_2^z$	$-0.22279\,\sigma_2^z$	$0.12533\,\sigma_1^z$	$0.09462\,\sigma_1^z$	$0.03964\,\sigma_3^z$
$0.04084\,\sigma_3^x\sigma_2^y\sigma_1^y\sigma_0^x$	$0.04532\,\sigma_3^x\sigma_2^y\sigma_1^y\sigma_0^x$	$-0.05100\,\sigma_3^x\sigma_2^x\sigma_1^y\sigma_0^y$	$-0.05711\,\sigma_3^x\sigma_2^x\sigma_1^y\sigma_0^y$	$0.03964\,\sigma_2^z$

In this section, we discussed the small case study of molecular hydrogen, observing that the optimal ordering is described by alternating between fully commuting sets of terms. We proceed to present two generalized ordering schemes derived from this heuristic, and analyze their performance across a variety of molecular systems.

4. Generalized Ordering Strategies

While analysis of the hydrogen molecule yielded interesting results regarding the spread of ordering strategies, it would be desirable to find an ordering strategy that is effective in the general case. For this, we require an analysis of a variety of systems, and ideally those of a chemically interesting size. While the hydrogen study does not provide this, the optimal ordering at equilibrium bond length does provide us with a potential starting point for our investigation.

4.1. Methods

To contrast our ordering schemes against other possible alternatives, it is necessary to briefly review the conventions we use to describe other ordering schemes, previously discussed in Reference [34].

Perhaps the most immediately obvious ordering scheme is the *magnitude* ordering. Here, terms are ordered according to the magnitude of their coefficient, from largest to smallest—the term with the largest magnitude coefficient is simulated first, followed by the second largest magnitude coefficient, and so on. This ordering scheme has the immediate appeal that high magnitude terms are simulated first. As a result, one could expect a reduction in the number of high magnitude terms in the error operator. Loosely, this is due to implementing high magnitude terms first, as opposed to later in the sequence where they may compound earlier errors. However, this approach does not take into consideration the structure of the error operator—high magnitude terms may not result in a meaningful increase in error if, for example, they commute with many other terms.

The *lexicographic* ordering is an ordering scheme which attempts to maximize the similarity of the Pauli strings of adjacent terms. This is essentially a numerical ordering with respect to the Pauli strings. While this scheme may seem arbitrary, it is known to result in a maximum amount of gate cancellation, with commensurate reduction in overall quantum computational cost [33]. However, there is little reason to suspect that this ordering would be beneficial for the purposes of Trotter error. Such an increase could lead to the requirement for a greater number of Trotter steps, undermining its advantages. It should be noted that our strategy for ordering terms lexicographically is a small modification from that used in other work [33]. Whereas other approaches have grouped the qubit Hamiltonian terms by their fermionic operator, our approach does not maintain this structure, and instead stores the qubit Hamiltonian as a list of individual weighted Pauli string terms. As such, we do not treat any terms based on their fermionic role. Similarly, our ordering is based purely on the structure of the individual Pauli strings, rather than the fermionic terms they correspond to.

Testing on random Hamiltonians [49] proved relatively optimistic for the prospects of some of the ordering schemes described in later sections of this paper. However, it has been established that testing on random Hamiltonians does not realistically capture the behavior of real molecular Hamiltonians when considering Trotter errors. A more rigorous analysis requires the use of real chemical Hamiltonians. We therefore used a set of 44 molecular Hamiltonians, largely using the same dataset as in Reference [34]. Table 2 shows a breakdown of these systems. Equilibrium molecular geometries were gathered from the NIST CCCBDB database optimized at the Hartree–Fock level [50]. Molecular orbital integrals in the Hartree–Fock basis were obtained from Psi4 [47] and OpenFermion [48]. As in Reference [34], the procedure for determining the exact Trotter error was to directly calculate the expectation value of the Trotterized Hamiltonian with the exact ground state, using our Python code. The Trotter evolution time for each Hamiltonian was set

$$
t = \begin{cases} 1, & \text{where } |E_{FCI}| < 2\pi \\ \dfrac{1}{2\pi \left\lceil \frac{|E_{FCI}|}{2\pi} \right\rceil}, & \text{otherwise} \end{cases} \tag{7}
$$

using the full configuration interaction energy for computational convenience. In a real simulation, the energy provided by a lower level of theory would be used.

Table 2. The molecular dataset used. Note that most of the systems involving a non-minimal basis set were H_2 and HeH^+ systems, as specified in Appendix B. The polyatomic category includes molecules, ions and radicals.

Qubits	1–10	11–20	21–30	Total
Polyatomic	1	14	5	20
Atoms	4	6	0	10
Ions	2	2	0	4
Other bases	4	4	2	10
Total	11	26	7	44

4.2. Results—Magnitude Ordering

To provide some context for our discussion of Trotter ordering strategies, we first consider the Trotter error incurred for our systems using a magnitude ordering. Figure 4 demonstrates this against a variety of properties. The majority of the systems demonstrated Trotter error below the threshold for chemical accuracy. While there is little indication of increasing Trotter error with the number of spin-orbitals, it should be noted that the systems studied are all within the regime that can be practically simulated on modern classical computers. It is unclear as to whether this trend will persist beyond this regime. Nonetheless, this is encouraging for experimental studies in the near future, where the number of available qubits will be highly constrained.

There is no obvious correlation between the Trotter error and either the number of spin-orbitals or the number of terms in the Hamiltonian. However, when plotted against the maximum nuclear charge (i.e., the nuclear charge of the heaviest atom in the system)—as suggested by the previous work of Babbush et al. [36]—there is a notable, albeit loose, trend. All of the systems with single step magnitude ordering Trotter error insufficient for chemical accuracy involve only H and He nuclei. This could be affected by the special handling of the Trotter time for low energy systems in Equation (7). The results here reinforce the need for benchmarks of quantum approaches to electronic structure theory problems to consider systems with heavy atoms.

Beyond these systems, there is a peak in Trotter error around a maximum nuclear charge of 11. This is in agreement with the results of Babbush et al. [36], where it is observed that systems with mostly full spin-orbitals will incur low Trotter error. The peak in our data is due to the presence of Na or Mg, which in an STO-3G basis have many unfilled spin-orbitals.

4.3. Statistics of Commuting Hamiltonian Subsets

One approach to using commutativity structure to inform ordering strategy is to consider coloring of the incompatibility graph of the qubit Hamiltonian. As in Figure 2, we can consider this structure by representing the commutativity of the terms in the Hamiltonian as a graph, with terms in the Hamiltonian corresponding to nodes and an edge representing the case where terms do not commute. Generating such a graph requires $\mathcal{O}\left(N_{terms}^2\right)$ time, as calculating the commutator of two arbitrary terms is classically efficient.
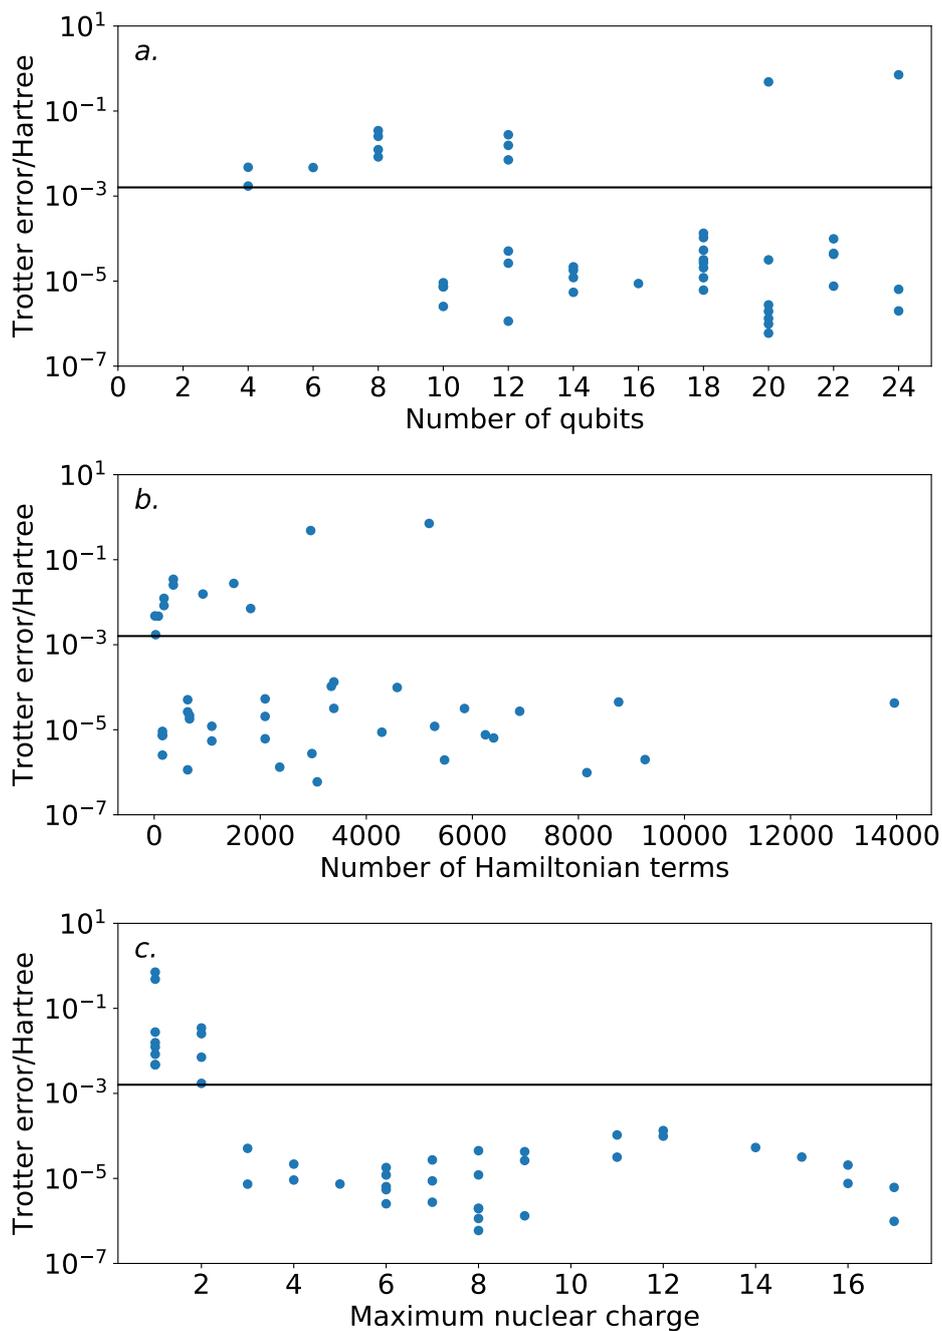
Figure 4. Trotter errors for the dataset of molecular Hamiltonians using a magnitude ordering. The vertical bar indicates chemical accuracy. Most of the systems achieve chemical accuracy with one Trotter step. (**a**): versus the number of spin-orbitals. Most of the high-error results are for low numbers of spin-orbitals. (**b**): versus the number of terms in the Hamiltonian. Again, most of the high-error results are for low numbers of terms. (**c**): versus the maximum nuclear charge. All of the high-error systems are for systems with exclusively light atoms, and the overall trend roughly follows the predictions of prior literature.

Following our example of the Hydrogen molecule, a strategy of dividing the Hamiltonian into sets of mutually commuting subsets—sets of terms where all members commute—can be followed. This is equivalent to finding a coloring of the incompatibility graph. Unfortunately, it is well known that finding a coloring with a minimal number of colors is NP-hard [51]. However, heuristics—often using a greedy approach—have been developed [52]. There is no immediately obvious reason to suspect that coloring the Hamiltonian using a heuristic is problematic for the purposes of our analysis. Indeed, these heuristics have seen frequent recent use for partitioning electronic structure Hamiltonians, for the purpose of measurement reduction in variational quantum algorithms [53–56]. Nonetheless, it does suggest that caution should be used when considering the generalizability of the results presented.

Graphs representing the Hamiltonians in our dataset were generated using the NetworkX Python package [57]. Due to the relative computational ease of this procedure, our dataset here was extended by 16 additional systems, as indicated by Appendix B. From here, colorings were generated using the greedy coloring method provided by the same package, using an independent set strategy [52]. Coloring schemes for both Jordan–Wigner and Bravyi–Kitaev Hamiltonians were generated. Figure 5 shows the number of independent sets found with regard to both the number of terms in the overall Hamiltonian, and the number of spin-orbitals describing the Hamiltonian. As is to be expected, the number of sets increases with the number of terms in the Hamiltonian.

There are $\mathcal{O}(N^4)$ terms in the initial electronic Hamiltonian. Of these, terms that do not share any indices with each other will clearly commute, even in the absence of the use of the Jordan–Wigner or Bravyi–Kitaev transformations. As such, it is reasonable to expect the number of independent sets to be $\mathcal{O}(N^3)$. Therefore the ratio of the number of terms and independent sets found can be expected to scale roughly linearly with the number of spin-orbitals involved. This scaling is shown in Figure 5, outside the smallest Hamiltonians. Notable outliers are present, likely due to molecular symmetries. Little difference is observed between the Jordan–Wigner and Bravyi–Kitaev Hamiltonians. This is to be expected, as the commutativity structure is defined by the physical electronic Hamiltonian; there is no obvious reason to expect that the mapping technique used would significantly affect this. Our results here are in agreement with those in Reference [56], where the same scaling was observed for the dual problem (i.e., coloring the compatibility graph of the Hamiltonian).

Figure 5 also demonstrates the average and standard deviation of the size of the independent sets found for each Hamiltonian. The increased average size of the groups is unsurprising, as the number of terms increases faster than the number of sets. The increasing variance in the size of the sets may be undesirable. This is due to the fact that evenly-sized groups could be advantageous for ordering schemes, due to an increased ability to distribute the placement of terms. It is possible that the use of an alternative coloring strategy could circumvent this, and further work is required to assess whether this has an impact on the ordering schemes presented in Section 4.4.

4.4. Subset-Based Ordering Schemes

The purpose here of dividing the Hamiltonian into mutually commuting subsets is to examine how a Trotter ordering which takes this into account will perform. The above analysis of the hydrogen molecule suggested that alternating between commuting sets may be an effective scheme for Trotter ordering. With Hamiltonians partitioned into commuting subsets as above, this scheme can be extended to larger systems. In this method, an ordered Hamiltonian is generated by sequentially picking terms from each subset until all subsets are depleted.

Two methodologies for this ordering algorithm were considered. In the first—the *depleteGroups* strategy—the sets were cycled through, picking the highest magnitude term from each and appending this to the ordered Hamiltonian. The second—the *equaliseGroups* strategy—at each stage picked the highest magnitude term in the largest subset, appending it to the ordered Hamiltonian. Where there were multiple "largest subsets", the highest magnitude term in the union of all largest subsets was used. Whereas the former strategy ensures that the sets are consistently cycled through until depletion, the latter ensures that the sets are evenly distributed throughout the Trotterized Hamiltonian.
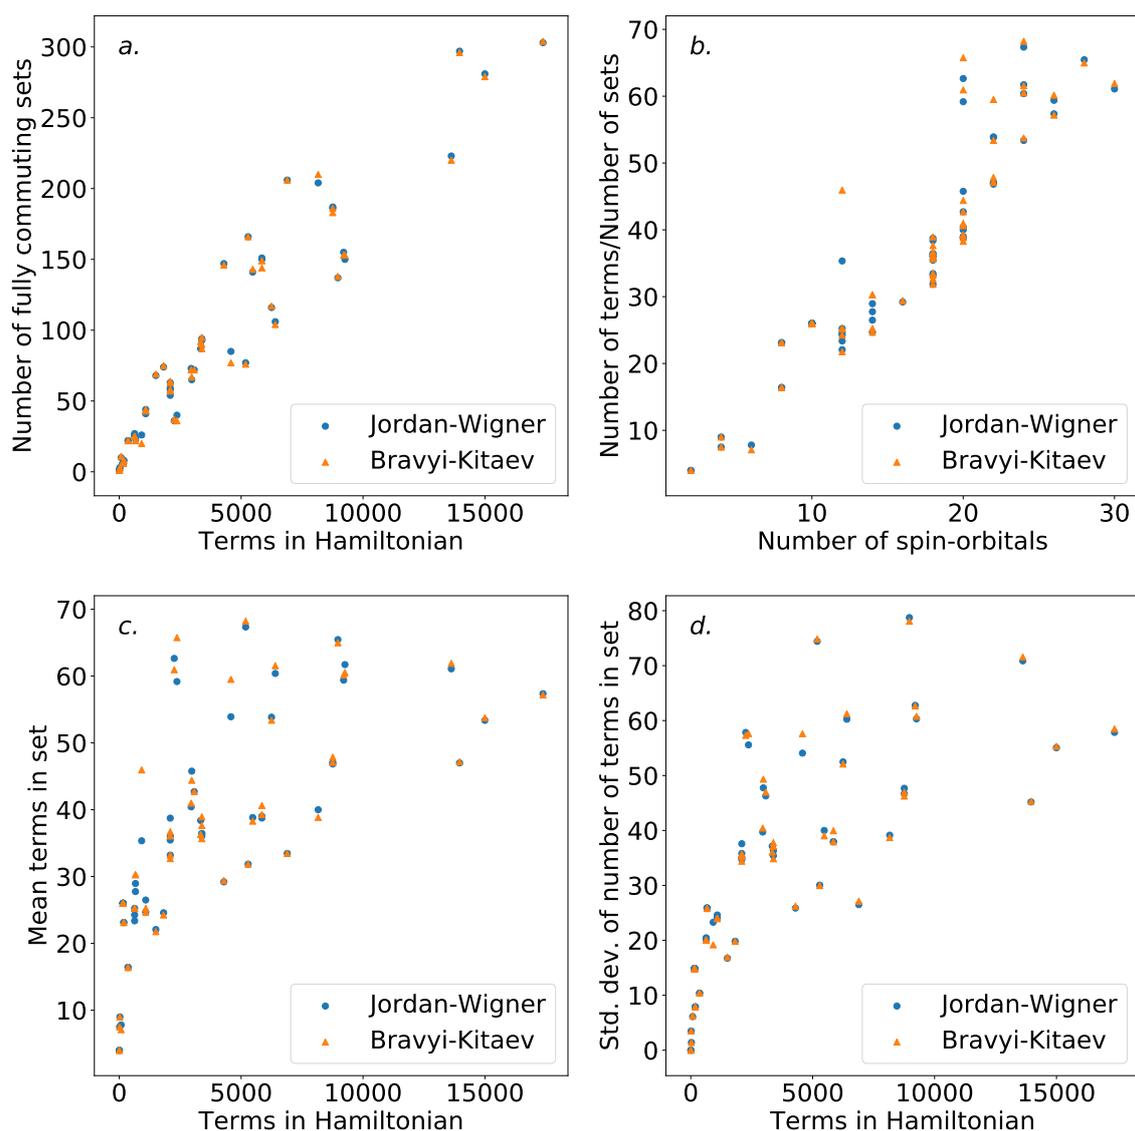
Figure 5. Statistics of the fully commuting sets of terms found in the coloring of the Hamiltonians in the dataset. (**a**): number of fully commuting sets versus the number of terms in the Hamiltonian. (**b**): number of independent sets divided by the number of terms, versus the number of spin-orbitals. A roughly linear trend is observed, indicating a $\Theta\left(N^3\right)$ scaling. (**c**): average number of terms in each fully commuting subset for a given Hamiltonian. (**d**): standard deviation of number of terms in each fully commuting subset for a given Hamiltonian. The increasing variance in group sizes could be problematic for ordering purposes.

The Trotter error for one Trotter step (using a second-order Trotter approximation) was calculated for each Hamiltonian using both ordering strategies, using the approach discussed in Section 4. Clearly, as this includes systems requiring more than 8 qubits, an exhaustive search of all possible orderings is not possible for all but the smallest of the systems included. Instead, we compare the depleteGroups and equaliseGroups orderings against other ordering schemes described above. We consider the performance of each ordering in terms of the number of spin-orbitals considered in the molecular system, and in the maximum nuclear charge.

Figure 6 shows the results of this analysis. High variance in ordering performance is associated with systems with only light atoms, but for most systems with heavier atoms, there is only a roughly 0.0001 a.u. difference in Trotter error between the orderings. This is a large variance in comparison to the absolute Trotter error. However, Figure 4 shows that using a magnitude ordering for these systems yielded an error that is sufficient for chemical accuracy. The observed difference between orderings is insufficient to increase the error above this threshold. As such, these results suggest that ordering choice when performing VQE can be determined by other factors, such as circuit length minimization. This is in sharp contrast to the results in Section 3 on the hydrogen molecule in an STO-3G basis, where the Trotter error was dramatically impacted by the Trotter ordering, resulting in a differing amount of Trotter steps required for chemical accuracy. This reinforces the need to consider larger systems when benchmarking methods. Conversely, if phase estimation is to be performed, it is likely that this error would be compounded due to applications of higher powers of the Trotterized unitary. In this context, an optimal ordering scheme in terms of Trotter error would be more necessary.

The equaliseGroups and depleteGroups strategies appear reasonably promising. The depleteGroups scheme is superior to the equaliseGroups strategy in all bar one of the systems with more than 20 spin-orbitals, and performs roughly commensurately elsewhere. We conclude that of the depleteGroups and equaliseGroups strategies, the depleteGroups strategy should be favored, although it is possible that this is a result of the uneven group sizes shown by Figure 5. An inspection of Figure 6 shows that the depleteGroups strategy is better than the magnitude ordering in eighteen cases.

More consistency is observed with respect to the maximum nuclear charge. Highly variable performance is observed with light atoms; however beyond this, the difference in Trotter error is relatively minor. For several of the systems consisting of the most (22) spin-orbitals, the depleteGroups strategy begins to dramatically outperform all other orderings considered—including the magnitude ordering. For systems involving period three atoms, there is only one exception to this result. More data are required to assess whether this trend consistently extends to other, and larger, systems.

It should be noted that the process reported here to color the Hamiltonian commutativity graph was relatively simple, examining only one possible ordering strategy using a standard library. Future work could investigate the impact of altering the details of this scheme. A variety of factors could be considered here. For example, it is not immediately obvious whether it would be preferable to split the graph into few large independent sets, or many smaller ones. As mentioned above, various schemes of Hamiltonian term partitioning have been developed to reduce the cost of variational quantum algorithms, by combining sets of commuting or anticommuting terms [53–56]. It remains an open question as to whether Trotter ordering schemes based on graph coloring heuristics can be applied to Hamiltonians with terms combined in such a fashion.

The results of the depleteGroups strategy are encouraging, although they do not achieve improvement over a magnitude ordering in all cases. Further work examining larger systems is required to test whether the observed improvement with the depleteGroups strategy is maintained at larger numbers of spin-orbitals, although this may prove computationally difficult.

In complement to the above discussion, an entirely different approach to ordering could be considered, which relies upon the Trotter error operator. As the norm of this can be efficiently (albeit slowly) determined classically, it is possible that this information could be used to directly inform an ordering strategy—rather than indirectly through the commutation relations of terms. For completeness, we consider such an approach now.
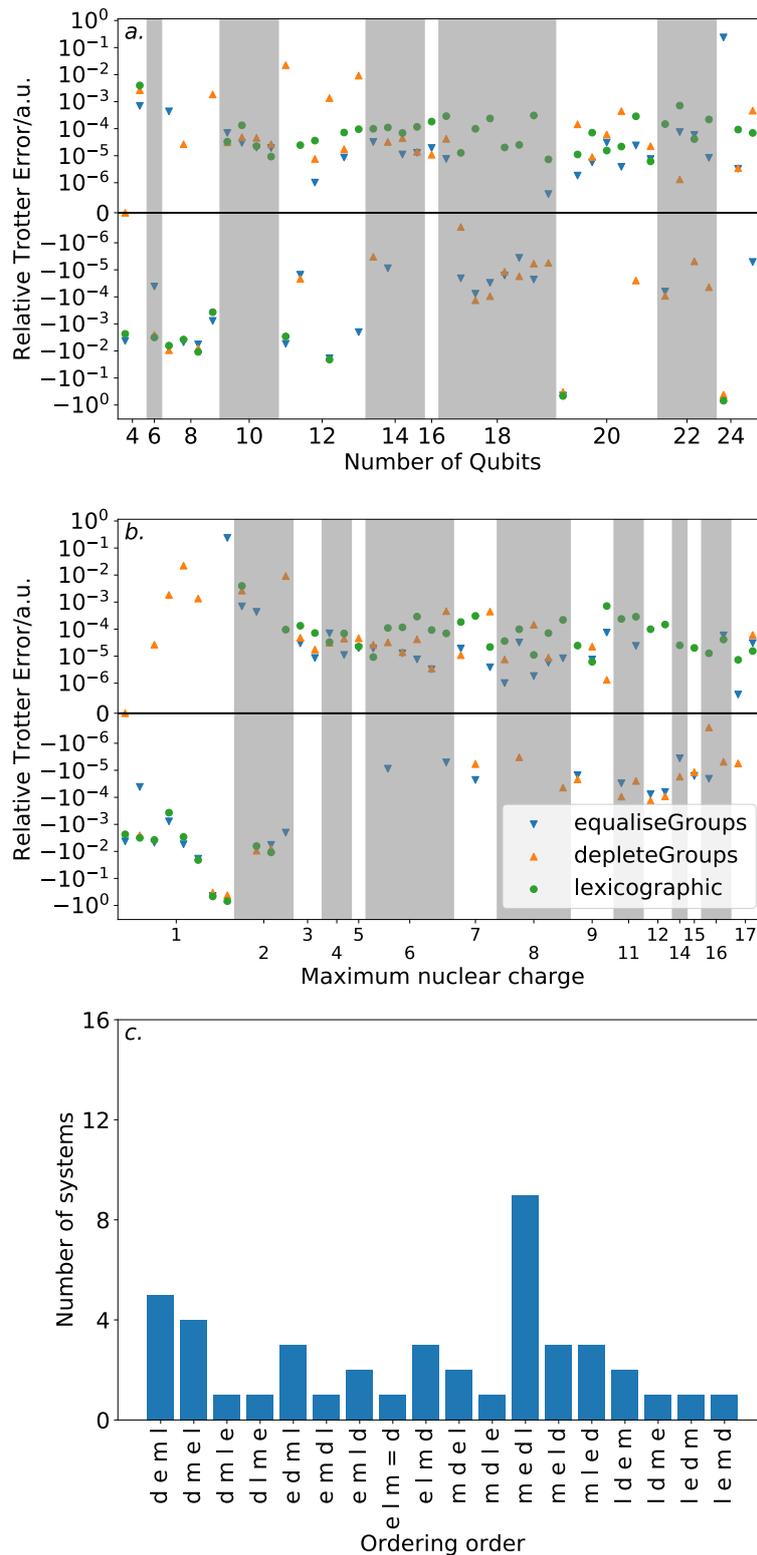
Figure 6. Trotter error of the depleteGroups and equaliseGroups orderings relative to a magnitude ordering. Upper plots are linear within $\pm 10^{-6}$. (**a**): by number of qubits. (**b**): by maximum nuclear charge. Again, the magnitude ordering is preferable in most cases; however, for systems with period three atoms, the depleteGroups and equaliseGroups are best. (**c**): frequency of "ordering orders", being the sequence of ordering performance. The distribution here is relatively flat.

5. Error Operator Strategies

5.1. Small Systems

As discussed above, the Trotter error operator can be used to determine the Trotter error of a given Trotterization scheme, with the norm of this operator giving a (loose) upper bound on the Trotter error. This information could be used to guide an ordering strategy.

To this end, we first examine how the Trotter error operator norm varies with respect to the Trotter ordering. The Trotter error operator norm of each ordering for the hydrogen molecule in a minimal basis was calculated, using a second-order Trotter approximation, with the fully commuting terms removed as in Section 3. Similarly, the Trotter error operator norm for 100,000 randomly chosen orderings of the helium hydride ion—also in a minimal basis—was calculated, along with the expectation of the Trotterized unitary with the exact ground state for each. As there are $27! \approx 10^{28}$ possible orderings for this Hamiltonian, our sample is only a small fraction of the entire ordering space.

The results of this analysis are shown in Figure 7. The two systems display different trends. For the hydrogen molecule, there is a loose correlation between the Trotter error operator norm and the the true Trotter error. The worst possible orderings—incurring an approximate factor of 5 increase in Trotter error versus the best orderings—also obtain a low error operator norm. This is encouraging, as it suggests that relying on the Trotter error operator norm will not result in an extremely poor ordering, at least for the hydrogen molecule in a STO-3G basis. However, it should be noted that within the low Trotter error region (below 0.001 a.u.), there is a broad spread of error operator norms, indicating that this approach may not be effective for finding truly optimal orderings. The helium hydride results are less clear. Indeed, the worst orderings tested in terms of true Trotter error have a relatively low error operator norm. Nonetheless, in both cases, the lowest true Trotter error orderings also resulted in a low Trotter error operator norm. As for the hydrogen molecule shown here and discussed above, the helium hydride results show a high density of orderings in the low Trotter error region—suggesting that most ordering schemes will be "good enough ", provided they do not systematically result in falling into the high Trotter error region.

5.2. Term Insertion Error Operator

While the Trotter error operator can be determined in polynomial time on a classical computer, the triple sum of triple commutators present in Equation (4) nonetheless presents some computational difficulty. An examination of the Trotter error operator suggests that it can be calculated in $\mathcal{O}(N_{terms}^3)$ time. Assuming that there are $\mathcal{O}(N_o^4)$ terms in the Hamiltonian (where N_o is the number of spin-orbitals involved), the calculation of each Trotter error operator requires $\mathcal{O}(N_o^{12})$ time. Where the Trotter error operator is calculated a constant number of times, this is surmountable for low numbers of spin-orbitals, due to the sparsity of the coefficients in the electronic Hamiltonian. However, an ordering scheme based on the calculation of the Trotter error operator may require repeated calculation of varying Hamiltonians. If the error operator must be evaluated an average of $\frac{N_{terms}}{2}$ times for each of N_{terms} terms, the strategy would require time that scales roughly as the twentieth power of the number of orbitals. This is somewhat problematic.

In comparing the Trotter error operator of two orderings, it is therefore preferable to consider the difference between their error operators, thereby reducing the triple sum to a double sum. For clarity, we define the term operator

$$C_{\alpha\beta\gamma} = \frac{-\Delta_t^2}{12} \left[H_\alpha \left(1 - \frac{\delta_{\alpha,\beta}}{2} \right), [H_\beta, H_\gamma] \right] . \tag{8}$$
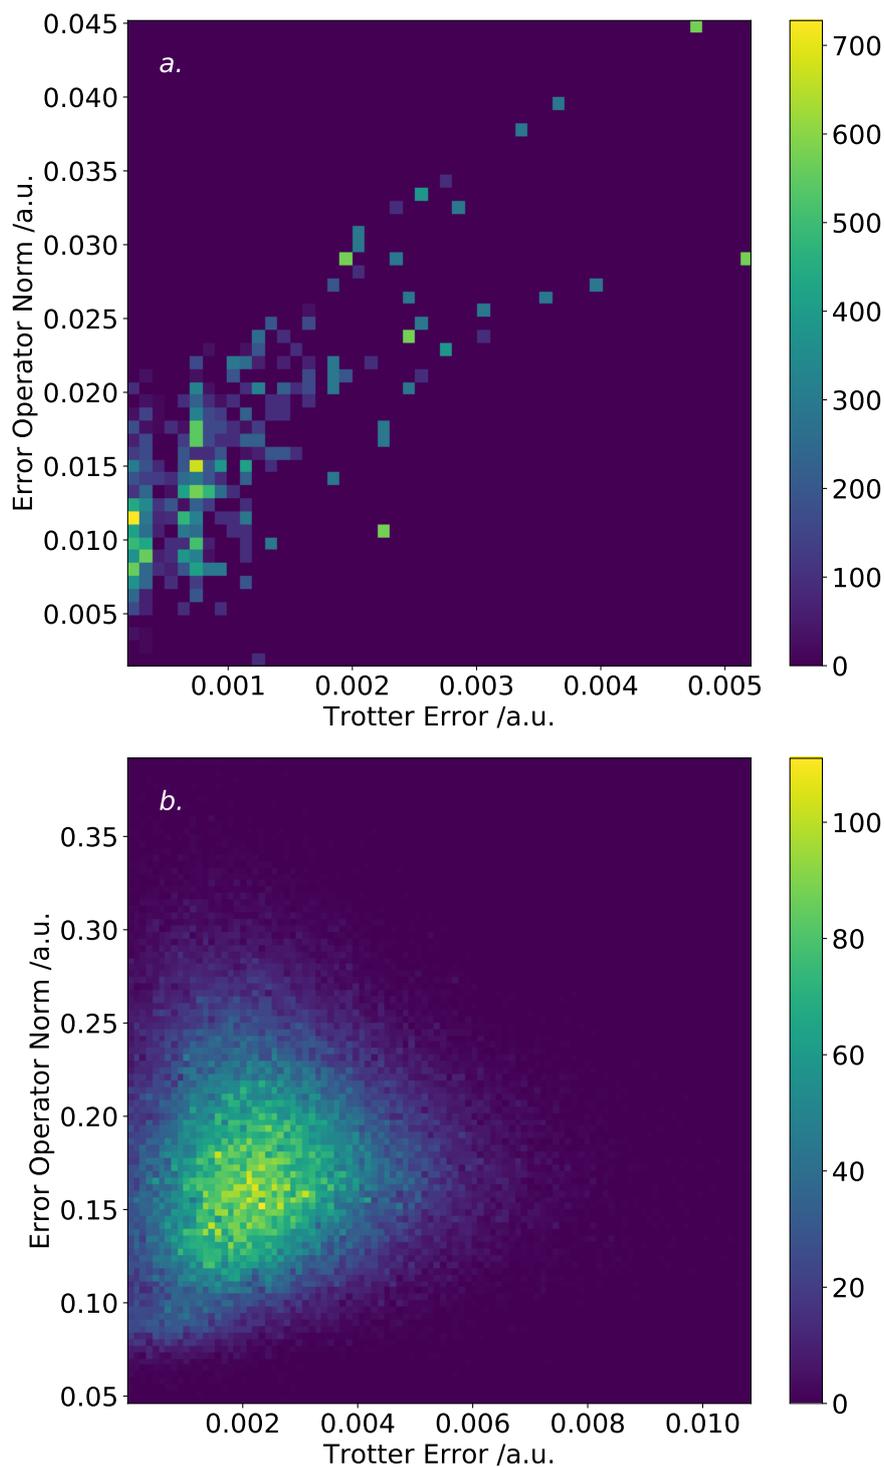
Figure 7. The Trotter error operator norm versus true Trotter error, for various orderings. (**a**): hydrogen molecule in a minimal basis. Two hundred and fifty bins are used. Loose correlation is observed, although there is ambiguity for a true Trotter error of less than 0.001 a.u. (**b**): helium hydride in a minimal basis, using 100,000 random orderings. One thousand bins are used. Little obvious correlation is observed.

If we insert a term H_i at index i into a given Trotterized Hamiltonian, and divide the β sum into $\beta < i$, $\beta = i$ and $\beta > i$ components, the new Trotter error operator V_i is then:

$$V_i = \sum_{\alpha \leq \beta} \sum_{\beta < i} \sum_{\gamma < \beta} C_{\alpha\beta\gamma} + \sum_{\alpha \leq i} \sum_{\gamma < i} C_{\alpha i \gamma} + \sum_{\alpha \leq \beta} \sum_{\beta > i} \sum_{\gamma < \beta} C_{\alpha\beta\gamma}. \tag{9}$$

We separate the terms where $\alpha = i$ or $\gamma = i$ from the final sum to yield:

$$V_i = \sum_{\alpha \leq \beta} \sum_{\beta < i} \sum_{\gamma < \beta} C_{\alpha\beta\gamma} + \sum_{\alpha \leq i} \sum_{\gamma < i} C_{\alpha i \gamma} + \sum_{\beta > i} \left(\sum_{\substack{\alpha \leq \beta \\ \alpha \neq i}} \sum_{\substack{\gamma < \beta \\ \gamma \neq i}} C_{\alpha\beta\gamma} + \sum_{\substack{\alpha \leq \beta \\ \alpha \neq i}} C_{\alpha\beta i} + \sum_{\gamma < \beta} C_{i\beta\gamma} \right) \tag{10}$$

The difference between V_i and V—the terms added to the error operator by the insertion of term H_i at index i—will be the terms that contain either $C_{i\beta\gamma}$, $C_{\alpha i \gamma}$ or $C_{\alpha\beta i}$. This is true of all terms in the single and double sums in Equation (10), and none of the terms in the triple sums. Replacing $C_{\alpha\beta\gamma}$ we therefore have

$$V_i' = \frac{-\Delta_t^2}{12} \left(\sum_{\alpha \leq i} \sum_{\gamma < i} \left[H_\alpha \left(1 - \frac{\delta_{\alpha,i}}{2} \right), [H_i, H_\gamma] \right] \right.$$

$$\left. + \sum_{\beta > i} \left(\sum_{\gamma < \beta} [H_i, [H_\beta, H_\gamma]] + \sum_{\substack{\alpha \leq \beta \\ \alpha \neq i}} \left[H_\alpha \left(1 - \frac{\delta_{\alpha,\beta}}{2} \right), [H_\beta, H_i] \right] \right) \right) \tag{11}$$

the Trotter term insertion error operator, which consists of the terms added to the error operator upon inserting a term H_i at index i. Provided a reference error operator, this is substantially easier to compute than the calculation of the entire error operator afresh.

5.3. Error Operator Based Ordering Schemes

We can now define a Trotter error operator based ordering scheme, following the logic of an insertion sort. We begin with the unsorted Hamiltonian. The highest magnitude term is then removed from the unsorted Hamiltonian, and forms the first term in the sorted Hamiltonian. The Trotter term insertion error operator (Equation (11)) is calculated in the cases where the second term is inserted either before or after the initial term. The norm of this operator is calculated in both cases. The scenario wherein the term insertion Trotter error operator norm is minimized is chosen as the correct branch, and the term is placed in the corresponding location. The third highest magnitude term is then tested in each of the three new potential locations and placed in the one which minimizes the term insertion Trotter error operator norm. This process is repeated until the unsorted Hamiltonian is depleted. The procedure for ordering the Hamiltonian in this manner is depicted as Figure 8.

Despite eliminating the need to calculate the Trotter error operator at every step in the optimization, the procedure remains extremely slow, requiring $\mathcal{O}\left(N_{terms}^4\right)$ time. As such, this ordering scheme was applied to a subset of 36 of the Hamiltonians in our dataset, and the one-step Trotter error calculated. Appendix B shows which systems are included in this subset.
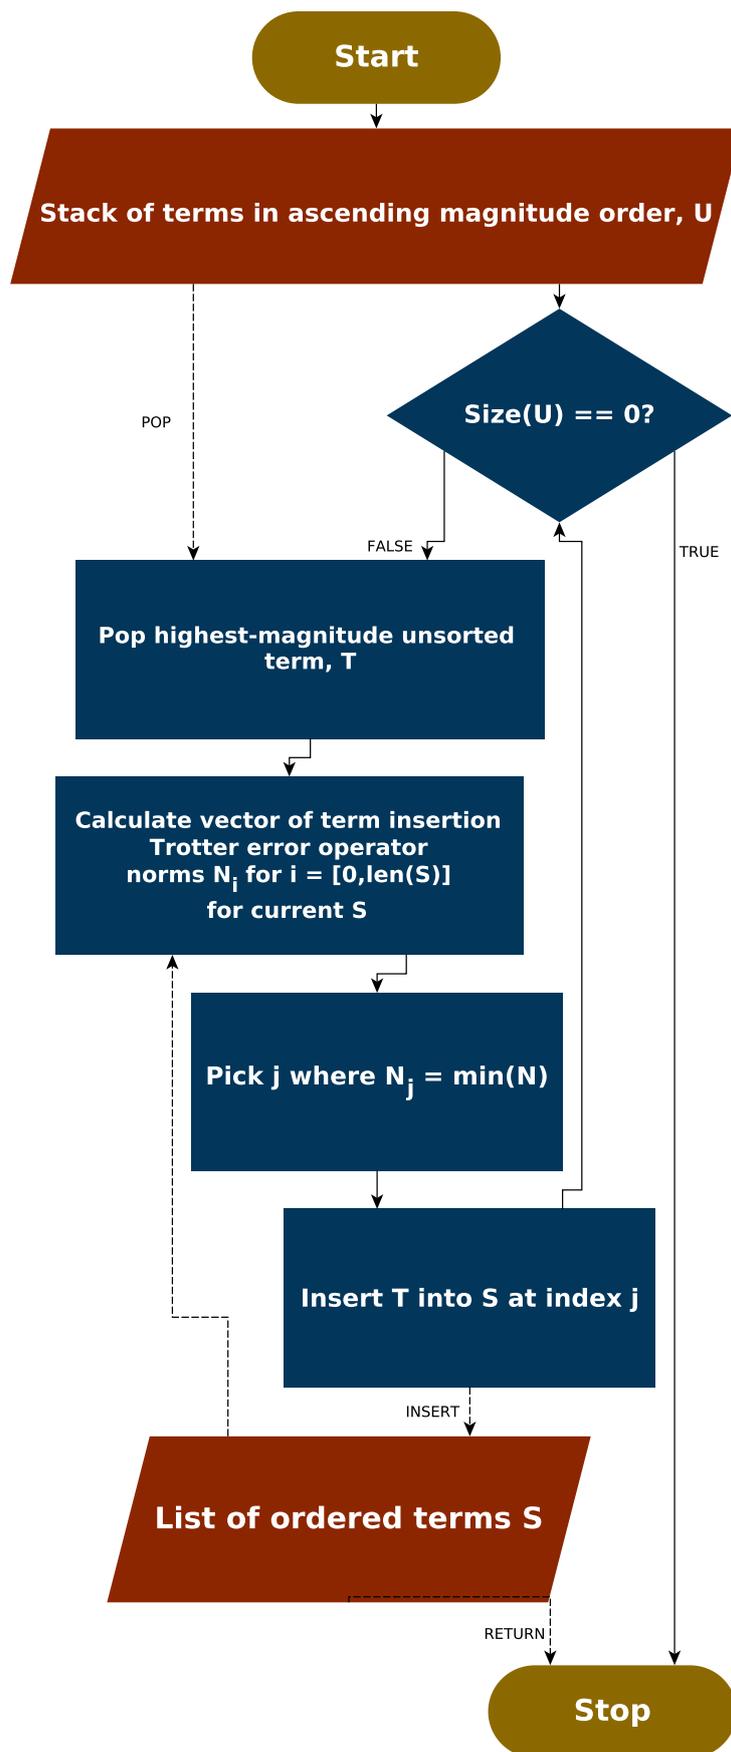
Figure 8. The procedure for performing the error operator norm minimization ordering.

Figure 9 shows the results of this analysis. Again, for large numbers of qubits or the inclusion of heavy atoms, the difference in Trotter error between the orderings is relatively minor. For systems with more than 12 qubits, the errorOperator ordering results in an approximately $\pm 10^{-4}$–$\pm 10^{-5}$ a.u. difference in error relative to a magnitude ordering. This is similar to that observed for the other ordering schemes considered. While outperforming the magnitude ordering in around half of the systems with maximum nuclear charge greater than 6, it does not consistently reach the performance of the depleteGroups ordering. Two possible inferences could be drawn from this. Firstly, the greedy optimization heuristic used to minimize the Trotter error operator norm may be at fault. This is particularly evidenced by the decreased performance of the errorOperator ordering for systems with 20 or more spin-orbitals, due to the large problem space. Alternatively, it could be simply due to the Trotter error operator norm being a loose bound on the actual Trotter error. While it is regrettable that this ordering does not appear to systematically outperform the other orderings considered, these factors—coupled with the computational difficulty of repeated calculations of the Trotter error operator—suggest that it will be difficult to find an effective and scalable ordering strategy reliant upon the Trotter error operator.

In several systems, most notably within those involving 18 spin-orbitals, the depleteGroups and errorOperator strategies result in extremely similar errors—a similarity which is not observed for the other ordering strategies. This indicates some replicated behavior, although further work is required to determine what this may be.

Other work has defined a term importance metric by considering the impact of a given fermionic term on the Hartree–Fock ground state [33,35]. To prevent dependence on the accuracy of the Hartree–Fock state, we do not consider this approach here. However, given the performance of the error operator ordering compared to a magnitude ordering, it is likely that such an approach could yield improvement with regard to the Trotter error incurred. Future work could aim to assess whether this is indeed the case.

In this section, we present an ordering scheme based on minimizing the Trotter error operator norm. The approach did not yield consistent improvement upon the depleteGroups strategy, which—in addition to the high computational difficulty of the task—suggests that approaches to ordering by use of the Trotter error operator norm may prove difficult.
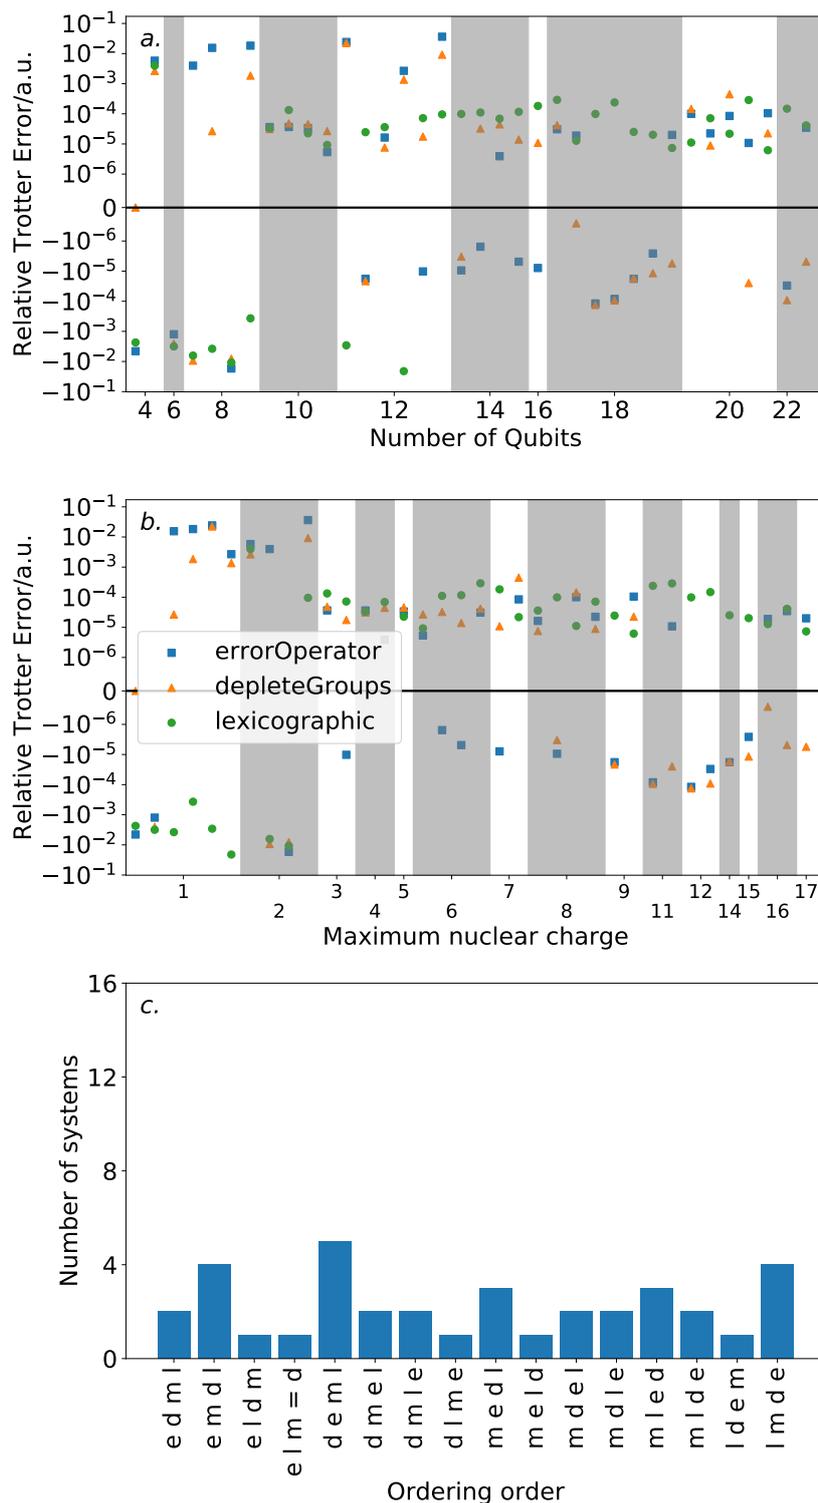
Figure 9. Trotter error of the errorOperator ordering relative to a magnitude ordering. Upper plots are linear between $\pm 10^{-6}$. (**a**): by number of qubits. (**b**): by maximum nuclear charge. As with the previous orderings, the variance between Trotter ordering schemes is low for systems involving heavy atoms. In these cases, the errorOperator ordering performs roughly commensurately with the magnitude ordering. (**c**): frequency of "ordering orders", being the sequence of ordering performance. The distribution here is relatively flat.

6. Discussion and Conclusions

In this paper, we discussed ordering schemes for Trotterization from a variety of perspectives. Brute-force analysis on the STO-3G molecular hydrogen Hamiltonian and inspection of the Trotter error operator show that the commutativity structure of the Hamiltonian is of vital importance in determining the Trotter error incurred by a given ordering. The commutativity structure—represented by the incompatibility graph—can be efficiently found on a classical computer. Greedy coloring heuristics can be used to partition this graph into commuting or anticommuting sets, a technique which has been used recently in the context of term reduction in variational quantum algorithms [53–56]. It remains an open question as to how these term reduction techniques can be related to Trotter ordering strategies when implementing variational quantum algorithms.

We reported two ordering schemes dependent on the term commutativity structure, which operate by partitioning the Hamiltonian into totally commuting sets of terms. Applying a greedy coloring heuristic results in a number of sets that scales approximately as $\mathcal{O}(n^3)$. By using these groups to inform placement of terms, this approach led to the best ordering scheme studied—the depleteGroups strategy. This strategy had the lowest Trotter error of all orderings in a plurality of Hamiltonians, although for systems involving heavy atoms, the difference when compared to a magnitude ordering was below the threshold for chemical accuracy.

Finally, we considered the use of the Trotter error operator norm to guide ordering strategies. A greedy algorithm aiming to minimize the norm of the Trotter error operator achieved performance consistent with magnitude ordering, though did not improve upon it in most cases. However, analysis of the Trotter error operator norm for very small molecules suggested that there is a high density of orderings which incur relatively low—if suboptimal—Trotter error.

Although the depleteGroups strategy showed promise, our analysis has indicated that finding an optimal Trotter ordering is a difficult task, due to the vast space of possible orderings. However, two results are of interest. Relative to the Trotter error, the difference between orderings can be substantial—in several cases, the difference in Trotter error incurred by the best and worst orderings was over two orders of magnitude. However, for the majority of systems studied, the overall Trotter error is extremely low. For all except four of the systems with more than 12 spin-orbitals, an ordering could be found that incurred a Trotter error below the threshold for chemical accuracy. Particularly for variational quantum algorithms—where Trotter decompositions are used to implement a variational Ansatz—this implies that Trotter error will not dominate the computational difficulty of simulations in the NISQ era.

Determining exact Trotter error on a classical computer is exponentially hard, due to the need to find an exact ground state. Therefore, there are two ways to scale examinations of Trotter error to systems involving larger numbers of spin-orbitals. The first is the use of efficient classical heuristics to compute *approximate* estimates of the Trotter error which outperform the Trotter error operator norm. In particular, the development of classical approximations which reproduce the relative performance of differing Trotter ordering strategies would be a useful direction for future work. The second—similar to other techniques using a quantum device to improve the quantum algorithm [58]—would be the use of NISQ devices themselves to directly evaluate ordering strategies through determination of exact Trotter errors in bulk. It is likely that a combination of each of these approaches will be necessary in order to determine ordering schemes that minimize Trotter error for large systems.

Author Contributions: Conceptualization, A.T., P.J.L., F.M., N.W. and P.V.C.; methodology, A.T., P.J.L., F.M., N.W. and P.V.C.; software, A.T.; data curation, A.T.; writing—original draft preparation, A.T.; writing—review and editing, A.T., P.J.L., F.M. and P.V.C.; supervision, P.J.L., F.M. and P.V.C.; project administration, P.J.L., F.M. and P.V.C.; and funding acquisition, P.J.L., F.M. and P.V.C.

Abbreviations

The following abbreviations are used in this manuscript:

CI   configuration interaction
FCI   full configuration interaction
VQE   variational quantum eigensolver
JW   Jordan–Wigner
BK   Bravyi–Kitaev

Appendix A. Alternative Ordering Schemes

Appendix A.1. Commutator Ordering

In this appendix, we discuss an additional two ordering schemes based on the commutativity structure of the Hamiltonian terms. We observe in Section 3 that the optimal ordering strategy for molecular hydrogen at equilibrium bond length was to divide the Hamiltonian into two mutually commuting subsets, and intersperse them so as to minimize the number of non-zero terms in the error operator, ordering each subset by magnitude. We present a general ordering scheme based on this approach in Section 4. However, generalizing this approach to larger systems is more difficult, as many decompositions of the Hamiltonian into mutually commuting subsets will be possible. For this reason, we now consider simpler ordering schemes.

This scheme—the *commutator* ordering—is still inspired by the optimal hydrogen ordering at equilibrium bond length. However, rather than attempting to find specific sets, it instead attempts to order by minimizing the amount of commutativity in the ordered Hamiltonian at each step.

We begin by ordering the Hamiltonian by coefficient magnitude. The highest magnitude term is then assigned to be the first term in the ordered Hamiltonian. We then consider the terms that do not commute with this, and the highest magnitude of these is appended to the ordered Hamiltonian. We keep track of the number of terms in the ordered Hamiltonian with which each term in the unordered Hamiltonian commutes, and we consider the subset of unordered Hamiltonian terms that minimize this number. Of these, the term with highest magnitude is chosen as the next term in the ordered Hamiltonian. This process is repeated until the unordered Hamiltonian is exhausted. Figure A1a demonstrates this process diagrammatically.

It should be noted that this process does not reproduce the exact optimal strategy for the molecular hydrogen test case, as in this instance the lower magnitude XY-set instead begins the ordering process. Nonetheless, it does produce an ordering which performs almost as well.

Appendix A.2. Reverse Commutator Ordering

In addition to this scheme, a slightly modified scheme was considered. In this scheme—the *reverseCommutator* ordering—we again start with an unordered Hamiltonian comprising of a list of terms. At each stage, instead of counting the number of terms in the ordered Hamiltonian that the candidate terms commute with, we count the number of terms in the *unordered* Hamiltonian that the candidate terms commute with. We then pick the terms that *maximize* this number, and append the highest magnitude of these. This process is depicted in Figure A1b.

Both algorithms are essentially greedy. The commutator ordering at each stage attempts to maximize the non-commutativity of the ordered Hamiltonian, whereas the reverseCommutator ordering at each stage attempts to minimize the non-commutativity of the unordered Hamiltonian.
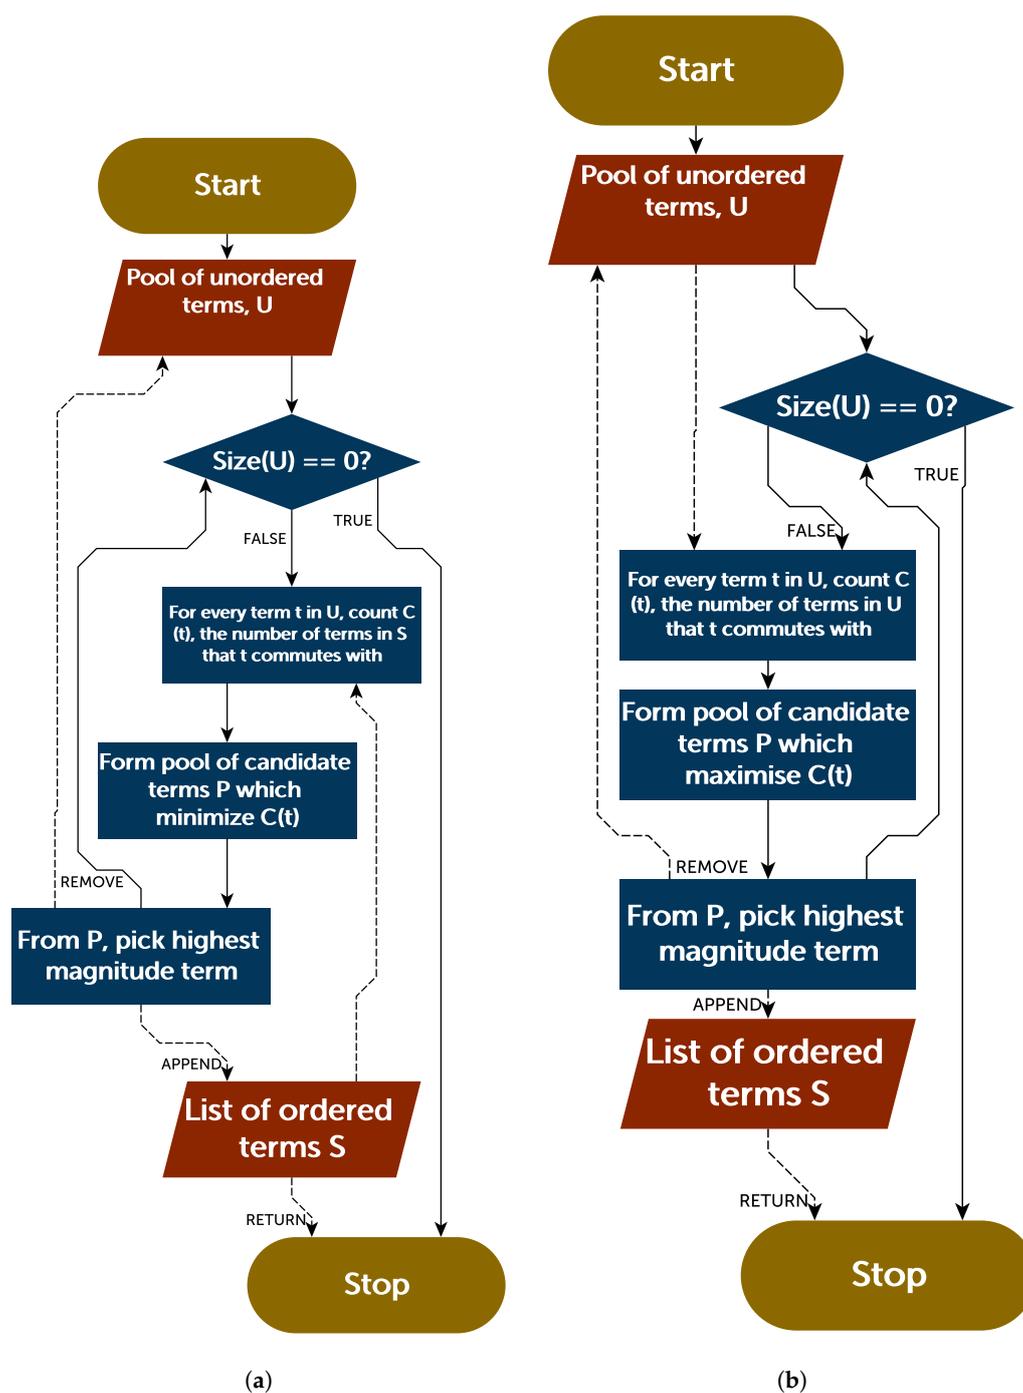
Figure A1. Flowcharts representing the commutator (**a**) and reverseCommutator (**b**) ordering schemes.

Appendix A.3. Performance of Commutator and ReverseCommutator Orderings

Figure A2 demonstrates the results of these simulations. We first note that, in most cases, the ordering schemes did not outperform a magnitude ordering. Versus the number of spin-orbitals, the commutator and reverseCommutator orderings do not fare well. While the reverseCommutator ordering outperformed other orderings by several orders of magnitude in some cases, the improvement is not consistent. However, assessing the orderings against the maximum nuclear charge indicates that in all bar two instances involving heavy atoms, the reverseCommutator ordering is the best of the three non-magnitude orderings considered. This suggests that, while the magnitude ordering (or the

depleteGroups ordering discussed in the main text) should be preferred for the reduction of Trotter error, there remains potential for ordering schemes based on commutativity structure.
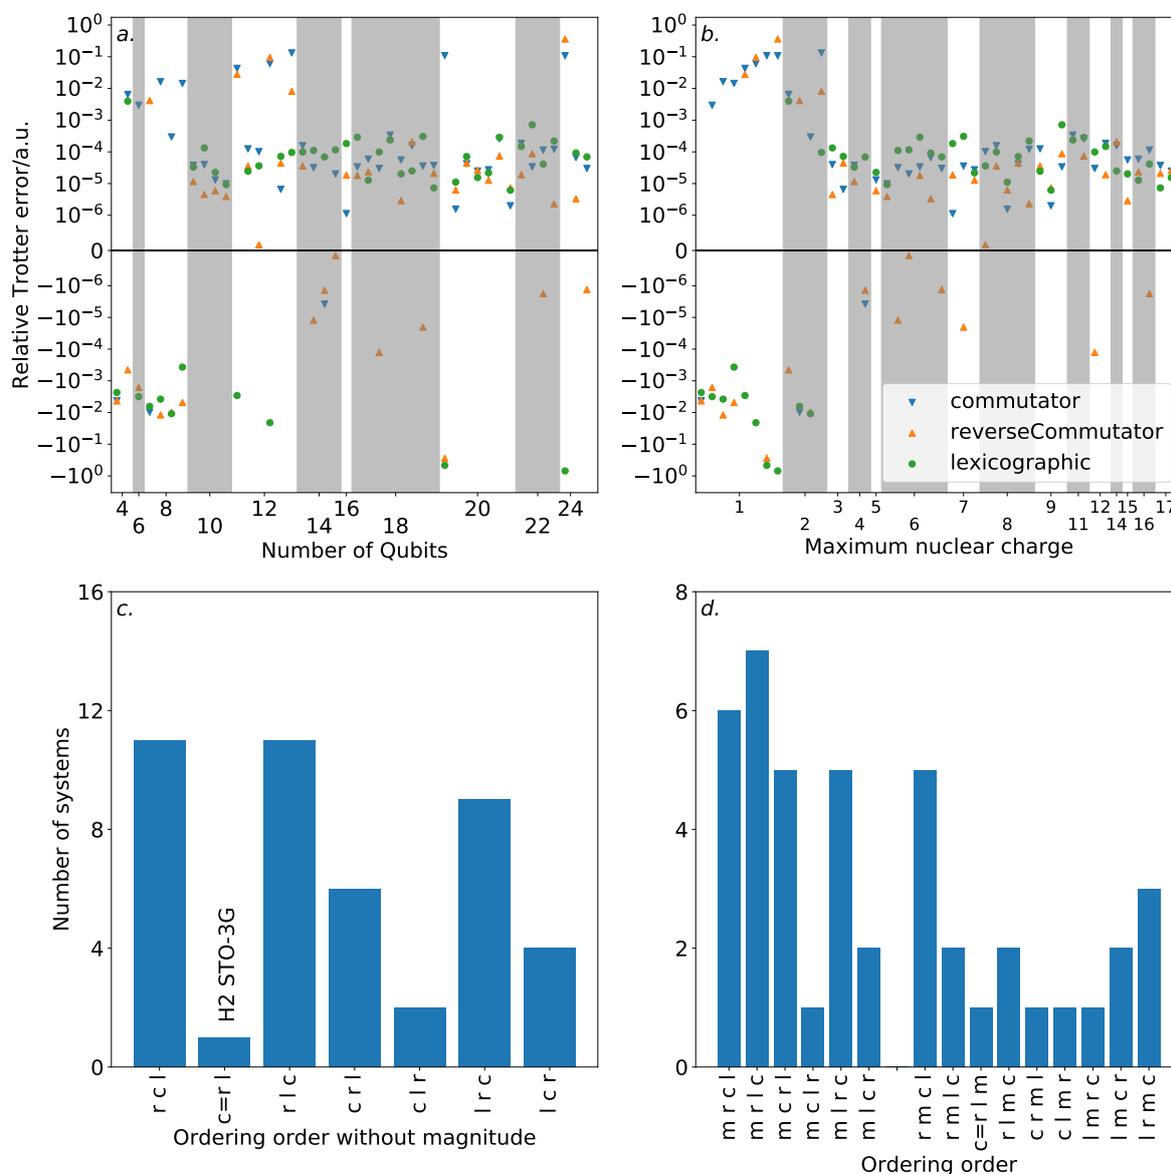

Figure A2. Trotter error of the commutator and reverseCommutator ordering relative to a magnitude ordering. Upper plots are linear within $\pm 10^{-6}$. (**a**): by number of qubits. (**b**): by maximum nuclear charge. The reverseCommutator ordering is best for almost all systems including heavy atoms, however in almost all cases, a simple magnitude ordering outperforms all orderings considered. (**c**): Frequency of "ordering orders", being the sequence of ordering performance, ignoring the magnitude ordering. (**d**): as lower left, but with the magnitude ordering.

The reliability of a magnitude ordering scheme does suggest improvement here. Our schemes used the magnitude of terms as only a tie-breaker in cases where candidate terms yielded a common level of non-commutativity against the ordered or unordered terms. The fact that the magnitude ordering yielded comparatively low error in most cases emphasizes the importance of the magnitude of the resultant terms in the error operator, rather than their mere presence. Our results here confirm the need for a more sophisticated effort at ordering selection, if minimization of Trotter error is desired. Additionally, the vast difference in ordering behavior between systems with light and heavy atoms reinforces the need to look at systems beyond STO-3G molecular hydrogen.

Appendix B. Systems in Dataset

Table A1. The systems included in the dataset. Geometries were obtained from the NIST CCBDB database [50], and molecular orbital integrals in the Hartee–Fock basis obtained from Psi4 [47] and OpenFermion [48].

System	Multiplicity	Charge	Basis	Qubits
B_1	2	0	STO-3G	10
Be_1	1	0	STO-3G	10
C_1O_1	1	0	STO-3G	20
C_1	3	0	STO-3G	10
Cl_1	2	0	STO-3G	18
F_2	1	0	STO-3G	20
H_1Cl_1 [1]	1	0	STO-3G	20
H_1F_1 [1]	1	0	3-21G	22
H_1F_1	1	0	STO-3G	12
H_1He_1	1	+1	3-21G	8
H_1He_1	1	+1	6-311G	12
H_1He_1	1	+1	6-31G	8
H_1He_1	1	+1	STO-3G	4
$H_1Li_1O_1$ [1]	1	0	STO-3G	22
H_1Li_1	1	0	STO-3G	12
H_1Na_1	1	0	STO-3G	20
H_1O_1	1	-1	STO-3G	12
H_2Be_1	1	0	STO-3G	14
$H_2C_1O_1$ [1]	1	0	STO-3G	24
H_2C_1	3	0	STO-3G	14
H_2C_1	3	0	STO-3G	14
H_2C_2 [1]	1	0	STO-3G	24
H_2Mg_1	1	0	STO-3G	22
H_2O_1	1	0	STO-3G	14
H_2S_1	1	0	STO-3G	22
H_2	1	0	3-21G	8
H_2 [1]	1	0	6-311G**	24
H_2	1	0	6-311G	12
H_2	1	0	6-31G**	20
H_2	1	0	6-31G	8
H_2	1	0	STO-3G	4
H_3N_1	1	0	STO-3G	16
H_3	1	+1	3-21G	12
H_3	1	+1	STO-3G	6
H_4C_1	1	0	STO-3G	18
H_4N_1 [1]	1	+1	STO-3G	18
Li_1	2	0	STO-3G	10
Mg_1	1	0	STO-3G	18
N_2	1	0	STO-3G	20
Na_1	2	0	STO-3G	18
O_2	3	0	STO-3G	20
P_1	4	0	STO-3G	18
S_1	3	0	STO-3G	18
Si_1	3	0	STO-3G	18

Table A2. The additional systems used for examining the statistics of fully commuting sets of terms in Section 4. Geometries were obtained from the NIST CCBDB database [50], and molecular orbital integrals in the Hartee–Fock basis obtained from Psi4 [47] and OpenFermion [48].

System	Multiplicity	Charge	Basis	Qubits
Ar_1	1	0	STO-3G	18
C_1O_2	1	0	STO-3G	30
Cl_1	1	-1	STO-3G	18
F_1	2	0	STO-3G	10
H_1He_1	1	+1	6-31G**	20
H_1Li_1	1	0	3-21G	22
H_1	2	0	STO-3G	2
H_2C_1	3	0	3-21G	26
H_2O_2	1	0	STO-3G	24
H_4C_2	1	0	STO-3G	28
He_1	1	0	STO-3G	2
K_1	2	0	STO-3G	26
N_1	4	0	STO-3G	10
Ne_1	1	0	STO-3G	10
O_1	3	0	STO-3G	10
O_2	1	0	STO-3G	20

References

1. Deglmann, P.; Schäfer, A.; Lennartz, C. Application of Quantum Calculations in the Chemical Industry—An Overview. *Int. J. Quantum Chem.* **2015**, *115*, 107–136. [CrossRef]
2. Kortagere, S.; Lill, M.; Kerrigan, J. Role of Computational Methods in Pharmaceutical Sciences. *Methods Mol. Biol.* **2012**, *929*, 21–48. [CrossRef] [PubMed]
3. Szabo, A.; Ostlund, N.S. In *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*; Dover Publications: Mineola, NY, USA, 1996; pp. 231–270.
4. Helgaker, T.; Jørgensen, P.; Olsen, J. *Molecular Electronic-Structure Theory*; Wiley: Chichester, NY, USA, 2000.
5. Gan, Z.; Grant, D.J.; Harrison, R.J.; Dixon, D.A. The Lowest Energy States of the Group-IIIA–Group-VA Heteronuclear Diatomics: BN, BP, AlN, and AlP from Full Configuration Interaction Calculations. *J. Chem. Phys.* **2006**. [CrossRef] [PubMed]
6. Pople, J.A.; Krishnan, R.; Schlegel, H.B.; Binkley, J.S. Electron Correlation Theories and Their Application to the Study of Simple Reaction Potential Surfaces. *Int. J. Quantum Chem.* **1978**, *14*, 545–560. [CrossRef]
7. Møller, C.; Plesset, M.S. Note on an Approximation Treatment for Many-Electron Systems. *Phys. Rev.* **1934**, *46*, 618. [CrossRef]
8. Aspuru-Guzik, A.; Dutoi, A.D.; Love, P.J.; Head-Gordon, M. Simulated Quantum Computation of Molecular Energies. *Science* **2005**, *309*, 1704–1707. [CrossRef]
9. Kassal, I.; Whitfield, J.D.; Perdomo-Ortiz, A.; Yung, M.H.; Aspuru-Guzik, A. Simulating Chemistry Using Quantum Computers. *Annu. Rev. Phys. Chem.* **2011**, *62*, 185–207. [CrossRef]
10. Cao, Y.; Romero, J.; Olson, J.P.; Degroote, M.; Johnson, P.D.; Kieferová, M.; Kivlichan, I.D.; Menke, T.; Peropadre, B.; Sawaya, N.P.D.; et al. Quantum Chemistry in the Age of Quantum Computing. *arXiv* **2018**, arXiv:1812.09976.
11. 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. *arXiv* **2017**, arXiv:1706.05413.
12. Lidar, D.A.; Wang, H. Calculating the Thermal Rate Constant with Exponential Speedup on a Quantum Computer. *Phys. Rev. E* **1999**, *59*, 2429–2438. [CrossRef]
13. Kassal, I.; Aspuru-Guzik, A. Quantum Algorithm for Molecular Properties and Geometry Optimization. *J. Chem. Phys.* **2009**, *131*, 224102. [CrossRef] [PubMed]

14. Reiher, M.; Wiebe, N.; Svore, K.M.; Wecker, D.; Troyer, M. Elucidating Reaction Mechanisms on Quantum Computers. *Proc. Natl. Acad. Sci. USA* **2017**, *114*, 7555–7560. [CrossRef] [PubMed]

15. 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. USA* **2008**, *105*, 18681–18686. [CrossRef] [PubMed]

16. 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.; et al. Towards Quantum Chemistry on a Quantum Computer. *Nat. Chem.* **2010**, *2*, 106–111. [CrossRef] [PubMed]

17. 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. [CrossRef]

18. 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.; et al. Scalable Quantum Simulation of Molecular Energies. *Phys. Rev. X* **2016**, *6*, 031007. [CrossRef]

19. 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. [CrossRef]

20. Hempel, C.; Maier, C.; Romero, J.; McClean, J.; Monz, T.; Shen, H.; Jurcevic, P.; Lanyon, B.P.; Love, P.; Babbush, R.; et al. Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator. *Phys. Rev. X* **2018**, *8*, 031022. [CrossRef]

21. Nam, Y.; Chen, J.S.; Pisenti, N.C.; Wright, K.; Delaney, C.; Maslov, D.; Brown, K.R.; Allen, S.; Amini, J.M.; Apisdorf, J.; et al. Ground-State Energy Estimation of the Water Molecule on a Trapped Ion Quantum Computer. *arXiv* **2019**, arXiv:1902.10171 .

22. 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. [CrossRef]

23. McClean, J.R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. The Theory of Variational Hybrid Quantum-Classical Algorithms. *New J. Phys.* **2016**, *18*, 023023. [CrossRef]

24. Smart, S.E.; Mazziotti, D.A. Quantum-Classical Hybrid Algorithm Using an Error-Mitigating N-Representability Condition to Compute the Mott Metal-Insulator Transition. *Phys. Rev. A* **2019**, *100*, 022517. [CrossRef]

25. Wei, S.; Li, H.; Long, G. Full Quantum Eigensolver for Quantum Chemistry Simulations. *arXiv* **2019**, arXiv:1908.07927.

26. Farhi, E.; Goldstone, J.; Gutmann, S. A Quantum Approximate Optimization Algorithm. *arXiv* **2014**, arXiv:1411.4028 .

27. Iliyasu, A.M.; Venegas-Andraca, S.E.; Yan, F.; Sayed, A. Hybrid Quantum-Classical Protocol for Storage and Retrieval of Discrete-Valued Information. *Entropy* **2014**, *16*, 3537–3551. [CrossRef]

28. Zhu, D.; Linke, N.M.; Benedetti, M.; Landsman, K.A.; Nguyen, N.H.; Alderete, C.H.; Perdomo-Ortiz, A.; Korda, N.; Garfoot, A.; Brecque, C.; et al. Training of Quantum Circuits on a Hybrid Quantum Computer. *Sci. Adv.* **2019**, *5*. [CrossRef]

29. Kitaev, A.Y. Quantum Measurements and the Abelian Stabilizer Problem. *arXiv* **1995**, arXiv:quant-ph/9511026.

30. Lloyd, S. Universal Quantum Simulators. *Science* **1996**, *273*, 1073–1078. [CrossRef]

31. Abrams, D.S.; Lloyd, S. Simulation of Many-Body Fermi Systems on a Universal Quantum Computer. *Phys. Rev. Lett.* **1997**, *79*, 2586–2589. [CrossRef]

32. Romero, J.; Babbush, R.; McClean, J.R.; Hempel, C.; Love, P.J.; Aspuru-Guzik, A. Strategies for Quantum Computing Molecular Energies Using the Unitary Coupled Cluster Ansatz. *Quantum Sci. Technol.* **2018**, *4*, 014008. [CrossRef]

33. Hastings, M.B.; Wecker, D.; Bauer, B.; Troyer, M. Improving Quantum Algorithms for Quantum Chemistry. *Quantum Inf. Comput.* **2015**, *15*, 1–21.

34. Tranter, A.; Love, P.J.; Mintert, F.; Coveney, P.V. A Comparison of the Bravyi–Kitaev and Jordan–Wigner Transformations for the Quantum Simulation of Quantum Chemistry. *J. Chem. Theory Comput.* **2018**, *14*, 5617–5630. [CrossRef] [PubMed]

35. 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.

36. Babbush, R.; McClean, J.; Wecker, D.; Aspuru-Guzik, A.; Wiebe, N. Chemical Basis of Trotter-Suzuki Errors in Quantum Chemistry Simulation. *Phys. Rev. A* **2015**, *91*, 022311. [CrossRef]

37. Jordan, P.; Wigner, E. Über Das Paulische Äquivalenzverbot. *Z. FÜR Phys.* **1928**, *47*, 631–651. (In German) [CrossRef]

38. Bravyi, S.B.; Kitaev, A.Y. Fermionic Quantum Computation. *Ann. Phys.* **2002**, *298*, 210–226. [CrossRef]

39. 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. [CrossRef]

40. 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. [CrossRef]

41. Setia, K.; Whitfield, J.D. Bravyi-Kitaev Superfast Simulation of Fermions on a Quantum Computer. *arXiv* **2017**, arXiv:1712.00446.

42. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*; Cambridge University Press: Cambridge, UK, 2010; pp. 221–226.

43. Lee, J.; Huggins, W.J.; Head-Gordon, M.; Whaley, K.B. Generalized Unitary Coupled Cluster Wave Functions for Quantum Computation. *J. Chem. Theory Comput.* **2019**, *15*, 311–324. [CrossRef]

44. Hatano, N.; Suzuki, M. Finding Exponential Product Formulas of Higher Orders. In *Quantum Annealing and Other Optimization Methods*; Springer: Berlin/Heidelberg, Germany, 2005; pp. 37–68.

45. 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. [CrossRef]

46. Kirby, W.M.; Love, P.J. Contextuality Test of the Nonclassicality of Variational Quantum Eigensolvers. *arXiv* **2019**, arXiv:1904.02260 .

47. 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.; et al. Psi4 1.1: An Open-Source Electronic Structure Program Emphasizing Automation, Advanced Libraries, and Interoperability. *J. Chem. Theory Comput.* **2017**, *13*, 3185–3197. [CrossRef] [PubMed]

48. McClean, J.R.; Kivlichan, I.D.; Sung, K.J.; Steiger, D.S.; Cao, Y.; Dai, C.; Fried, E.S.; Gidney, C.; Gimby, B.; Häner, T.; et al. OpenFermion: The Electronic Structure Package for Quantum Computers. *arXiv* **2017**, arXiv:1710.07629.

49. Tranter, A. Quantum Chemistry and Quantum Computers—Testing the Bravyi-Kitaev Mapping and Trotter Order Optimisations. Ph.D. Thesis, Imperial College London, London, UK, November 2018.

50. Johnson, R.D., III. NIST Computational Chemistry Comparison and Benchmark Database NIST Standard Reference Database Number 101 Release 18. Available online: https://cccbdb.nist.gov/ (accessed on 9 Februay 2016).

51. Garey, M.R.; Johnson, D.S.; Stockmeyer, L. Some Simplified NP-Complete Problems. In Proceedings of the Sixth Annual ACM Symposium on Theory of Computing, New York, NY, USA, 1974; pp. 47–63. [CrossRef]

52. Kosowski, A.; Manuszewski, K. Classical Coloring of Graphs. In *Contemporary Mathematics*; Kubale, M., Ed.; American Mathematical Society: Providence, RI, USA, 2004; Volume: 352, pp. 1–19. [CrossRef]

53. Verteletskyi, V.; Yen, T.C.; Izmaylov, A.F. Measurement Optimization in the Variational Quantum Eigensolver Using a Minimum Clique Cover. *arXiv* **2019**, arXiv:1907.03358.

54. Izmaylov, A.F.; Yen, T.C.; Lang, R.A.; Verteletskyi, V. Unitary Partitioning Approach to the Measurement Problem in the Variational Quantum Eigensolver Method. *arXiv* **2019**, arXiv:1907.09040.

55. Crawford, O.; van Straaten, B.; Wang, D.; Parks, T.; Campbell, E.; Brierley, S. Efficient Quantum Measurement of Pauli Operators. *arXiv* **2019**, arXiv:1908.06942.

56. Zhao, A.; Tranter, A.; Kirby, W.M.; Ung, S.F.; Miyake, A.; Love, P. Measurement Reduction in Variational Quantum Algorithms. *arXiv* **2019**, arXiv:1908.08067.

57.   Hagberg, A.; Swart, P.; Daniel, S.C. Exploring Network Structure, Dynamics, and Function Using NetworkX. In Proceedings of the 7th Python in Science Conference (SciPy 2008), Pasadena, CA, USA, 21 August 2008; pp. 11–15.

58.   Lu, D.; Li, K.; Li, J.; Katiyar, H.; Park, A.J.; Feng, G.; Xin, T.; Li, H.; Long, G.; Brodutch, A.; et al. Enhancing Quantum Control by Bootstrapping a Quantum Processor of 12 Qubits. *Npj Quantum Inf.* **2017**, *3*, 1–7. [CrossRef]