

Improved Fault-Tolerant Quantum Simulation of Condensed-Phase Correlated Electrons via Trotterization

Ian D. Kivlichan^{1,2}, Craig Gidney³, Dominic W. Berry⁴, Nathan Wiebe⁵, Jarrod McClean¹, Wei Sun⁶, Zhang Jiang¹, Nicholas Rubin¹, Austin Fowler³, Alán Aspuru-Guzik^{7,8}, Hartmut Neven¹, and Ryan Babbush¹

¹Google Research, Venice, CA 90291, USA

²Department of Physics, Harvard University, Cambridge, MA 02138, USA

³Google Research, Santa Barbara, CA 93117, USA

⁴Department of Physics and Astronomy, Macquarie University, Sydney, NSW 2113, Australia

⁵Institute for Nuclear Theory, University of Washington, Seattle, WA 98195, USA

⁶Google Research, Mountain View, CA 94043, USA

⁷Department of Chemistry, University of Toronto, Toronto, Ontario M5G 1Z8, Canada

⁸Department of Computer Science, University of Toronto, Toronto, Ontario M5G 1Z8, Canada

July 10, 2020

Recent work has deployed linear combinations of unitaries techniques to reduce the cost of fault-tolerant quantum simulations of correlated electron models. Here, we show that one can sometimes improve upon those results with optimized implementations of Trotter-Suzuki-based product formulas. We show that low-order Trotter methods perform surprisingly well when used with phase estimation to compute relative precision quantities (e.g. energies per unit cell), as is often the goal for condensed-phase systems. In this context, simulations of the Hubbard and plane-wave electronic structure models with $N < 10^5$ fermionic modes can be performed with roughly $\mathcal{O}(1)$ and $\mathcal{O}(N^2)$ T complexities. We perform numerics revealing tradeoffs between the error and gate complexity of a Trotter step; e.g., we show that split-operator techniques have less Trotter error than popular alternatives. By compiling to surface code fault-tolerant gates and assuming error rates of one part per thousand, we show that one can error-correct quantum simulations of interesting, classically intractable instances with a few hundred thousand physical qubits.

1 Introduction

The physics of interacting electrons in the presence of external fields predicts the properties of many materials as well as the dynamics of most chemical reactions. While the dynamics of many such systems appear intractable for classical computers, quantum computers were originally introduced as universal simulators capable of efficient quantum dynamical simulation [1, 2]. Later work [3] combined the quantum phase estimation algorithm [4] with quantum simulation of fermions [5] to show that quantum computers can efficiently prepare ground states and estimate ground state energies of these

Ian D. Kivlichan: Corresponding author: ian.kivlichan@gmail.com

Ryan Babbush: Corresponding author: ryanbabbush@gmail.com

fermionic systems whenever an initial state can be prepared with non-vanishing overlap on the ground state (as is often the case for systems of interest [6]).

Of particular interest has been quantum simulations of the molecular electronic structure problem [7]. The first quantum circuits for these simulations [8] had $\mathcal{O}(N^{10})$ scaling where N is the number of single-particle basis functions [9]. Most early work in this area focused on performing time evolution of Hamiltonians in a Gaussian basis, by means of Trotter-Suzuki-based methods [10, 11]. By tightening bounds on the Trotter error [9, 12–14] and optimizing the implementation of Trotter steps [15, 16] this cost was brought down to an empirically observed $\mathcal{O}(N^6)$ scaling, which would likely be even lower using the techniques of [17] and [18].

Recent work has further reduced these scalings via application of linear combinations of unitaries [19] methods such as Taylor series [20, 21] and qubitization [22] either in combination with quantum signal processing [23] or used directly in phase estimation [24–27]. Another important innovation in the quantum simulation context has been the use of plane wave, rather than Gaussian, bases [28, 29]. Using plane waves simplifies the quantum computation and also extends simulations to the condensed-phase. The best quantum algorithms for plane wave basis electronic structure achieve either $\tilde{\mathcal{O}}(N^2)$ gates with $\mathcal{O}(N \log N)$ space [30] or $\mathcal{O}(N^3)$ gates with $\mathcal{O}(N)$ space [26]. There are also many papers which focus on these simulations in first quantization [31–35], the most efficient of which [36] obtains $\mathcal{O}(\eta^{8/3} N^{1/3})$ gate complexity and $\mathcal{O}(\eta \log N)$ space complexity, where η is number of particles.

Here, we will also explore quantum simulations of the Hubbard model [37], one of the simplest models of interacting electrons. Despite this simplicity, the Hubbard model demonstrates a range of correlated electron behaviors, and is considered a candidate model of high-temperature superconductivity in cuprates [38, 39]. Work focusing on quantum computations of the Hubbard model includes [5, 26, 40–46]. Of these, [45, 46] achieve very close to $\mathcal{O}(N)$ scaling, but via techniques that appear less practical than the $\mathcal{O}(N^2)$ scaling achieved in [26] owing to worse constant factors. However the bounds in [46] also apply to low-order Trotter methods such as ones studied here.

Only a few papers have assessed the viability of these simulations within a fault-tolerant cost model [26, 27, 47, 48]. We focus on the surface code [49] using lattice surgery [50, 51] due to its low qubit requirement for logical operations, and high threshold error rate with experimentally accessible planar connectivities. The relevant quantities for the surface code are the number of T and Toffoli gates in the circuit, and the maximum number of logical qubits required in the algorithm. Implementing T and Toffoli gates within the surface code requires the consumption of magic states. Generating these magic states is far more expensive than any Clifford gate. The work of [26] employs linear combination of unitaries techniques to show that classically intractable instances of plane wave electronic structure and the Hubbard model can be error-corrected with fewer than a million physical qubits in a matter of hours, assuming 10^{-3} physical error rates. Other recent work has demonstrated that many-body localization in the Heisenberg model may be simulated with fewer than a billion T gates [52], or under ten million T gates on specialized graphs [53].

We focus on simulating simple solid-state materials such as lithium hydride, graphite, diamond, etc., as well as the uniform electron gas and the Hubbard model, both in classically difficult regimes. We focus on the latter two problems due to their scientific importance, the wide range of phases they demonstrate, and their history as benchmarks for classical simulations [39, 54–58]. Ultimately, we show that classically intractable instances of these simulations can be performed in the surface code with a few hundred thousand physical qubits in less than an hour.

We achieve these results using second-order Trotter formulas. These formulas are further optimized through improved quantum circuits for the Trotter steps introduced in [28] and [29], and then we combine those Trotter steps with phase estimation techniques from [59, 60]. We show how Hamming weight phasing [61] can be used to asymptotically improve the T complexity of the split-operator step from $\mathcal{O}(N^2 \log(1/\epsilon))$ to $\mathcal{O}(N^2 + N \log N \log(1/\epsilon))$, where ϵ is the rotation synthesis precision. These are thus the most T-efficient electronic structure Trotter steps in the literature. Using a bound on Trotter error that is similar to (but tighter than) a bound derived in [9], we perform extensive numerics requiring high-performance computing resources to bound the number of Trotter steps required.

Previous work on quantum simulation has analyzed computations which seek a fixed additive

(intensive) error in the energy, $\Delta E = \Theta(1)$, rather than a fixed relative (extensive) error in the energy, $\Delta E = \Theta(N)$. But unlike with small molecule quantum chemistry, for condensed-phase systems one is often interested in the energy per unit cell of a system; after all, absolute energies are meaningless in the thermodynamic limit. For instance, when computing response properties (e.g., the Young’s modulus of a solid) the allowable error is plainly extensive with system size.

Because the Bardeen-Cooper-Schrieffer theory of superconductivity predicts an intensive excitation gap [62], one would need to seek a fixed additive error in the energy to prepare superconducting ground states of the Hubbard model. However, since the quasiparticle density grows linearly with system size, the energy difference between superconducting and normal metal phases is extensive (were this not the case, then an infinitesimal fluctuation would destroy the superconducting state). Put differently, for sufficiently large systems, we expect that a single quasiparticle excitation will not affect most local observables (e.g., correlation functions), implying that size extensive error is acceptable in this context. In this paper, we take the perspective that for condensed-phase quantum simulations there are valid reasons for targeting either intensive or extensive errors, and we analyze both situations.

We show that low-order Trotter methods perform especially well when used with phase estimation to target an extensive error in the energy. Allowing the target error to grow with system size significantly ameliorates the poor scaling of low-order Trotter methods with respect to time and precision. In this context, we show that practical simulations of the Hubbard model and plane wave electronic structure models can be performed with roughly $\mathcal{O}(1)$ and $\mathcal{O}(N^2)$ T complexities until the models are so large (e.g., $N > 10^5$ orbitals) that only one Trotter step is required to obtain target precision, or other assumptions of our analysis break down.

In [Section 2](#), we introduce the Hubbard and plane wave electronic structure Hamiltonians, and briefly discuss the two Trotter step algorithms that we study—the fermionic swap network [29] and split-operator Trotter steps [28]—as well as how we compute their fault-tolerant costs. We review a technique for combining repeated arbitrary rotations by the same angle [61] which we use in both Trotter steps, and discuss several extensions, in [Appendix A](#). We review the fermionic swap network Trotter step of [29] in greater detail in [Appendix B](#) and the split-operator Trotter step of [28] in [Appendix C](#), including several improvements to the algorithms. In [Appendix C.2](#), in particular, we introduce a scheme for asymptotically more efficiently simulating the potential energy operator in the split-operator Trotter step.

We then proceed to compute upper bounds on the fault-tolerant costs of accurately simulating the ground state energy of the uniform electron gas and the Hubbard model using phase estimation. We expand on our discussion of the fault-tolerant costs of each Trotter step in terms of the number of arbitrary rotations, Toffoli, and T gates from [Section 2](#) in [Appendix D](#), including upper bounds on the numbers of arbitrary rotations and T/Toffoli gates. We present an improved bound on the second-order Trotter error, and on the ground state energy shift, in [Appendix E](#); this bound does not require the product of the evolution time and norm of the Hamiltonian to be small.

In [Appendix F](#), we construct circuit primitives for phase estimation from Trotter steps for the electronic structure problem. In [Appendix G](#), we numerically evaluate the commutators which contribute to the Trotter error, allowing us to upper bound the shift in the ground state energy using the results of [Appendix E](#). We then combine these results in [Section 3](#) to compute upper bounds on the fault-tolerant costs of simulation. We use our Trotter analysis to place a numerical upper bound on the number of T gates required for fault-tolerant phase estimation for a range of system sizes. Together with the maximum number of logical qubits required, this allows us to determine the number of surface code physical qubits needed for distilling T states, the number of physical qubits to store and route data, and the execution time of these algorithms. Additional details on the precision targets we use, further numerical data, as well as some nuances of the minimization procedure, are detailed in [Appendix H](#). Finally, we discuss and summarize our results in [Section 4](#). We summarize the flow of the paper and how the different sections play into the broader whole in [Algorithm 1](#).

Algorithm 1: Determine fault-tolerant costs of estimating electronic structure problem ground state energies

Data: System parameters (electronic Hamiltonian, number of electrons η ; Section 2)
 Map physical Hamiltonian to qubit Hamiltonian using Jordan-Wigner transformation (Section 2);

Begin Phase estimation cost optimizations

- Use split-operator or fermionic swap network Trotter step for evolution within phase estimation (Section 2);
- Optimize Trotter step circuits to maximize simultaneous rotations for Hamming weight phasing (Section 2.1);
- Bound Trotter errors and maximum shift in eigenphases (Section 3.1);
- Determine precision target—either relative to total energy, or absolute (Section 3.2);
- Numerically minimize number of costly fault-tolerant gates (Section 3.2);
- Estimate surface code resource requirements (Section 3.3);

End

Result: Fault-tolerant costs of estimating system ground state energy (Section 3.3, Section 4)

2 Hamiltonians and Trotter Steps

This paper will analyze and reduce the cost of several approaches to simulating fermionic Hamiltonians of the form

$$H = \sum_{pq} T_{pq} a_p^\dagger a_q + \sum_p U_p n_p + \sum_{p \neq q} V_{pq} n_p n_q, \quad (1)$$

where a_p^\dagger and a_p are fermionic creation and annihilation operators and $n_p = a_p^\dagger a_p$ is the number operator for the corresponding spin-orbital. Mapping to qubits under the Jordan-Wigner transformation [63, 64], Eq. (1) becomes

$$\begin{aligned} H = & \sum_{p < q} \frac{T_{pq}}{2} (X_p Z_{p+1} \cdots Z_{q-1} X_q + Y_p Z_{p+1} \cdots Z_{q-1} Y_q) + \sum_{p < q} \frac{V_{pq}}{4} Z_p Z_q \\ & - \sum_p \left(\frac{T_{pp} + U_p}{2} + \sum_q \frac{(1 - \delta_{pq}) V_{pq}}{2} \right) Z_p + \sum_p \left(\frac{T_{pp} + U_p}{2} + \sum_q \frac{(1 - \delta_{pq}) V_{pq}}{4} \right) \mathbb{1}. \end{aligned} \quad (2)$$

This form includes a range of Hamiltonians. Note that there has also been a large body of work on using different fermionic encodings [65–71]; however, these methods do not appear to offer significant advantages in terms of the T complexity of simulations because they do not reduce the number of rotations that must be simulated.

First, we consider this Hamiltonian when the coefficients correspond to the plane wave electronic structure Hamiltonian. In the plane wave dual basis [28], the coefficients of Eq. (1) are

$$\begin{aligned} T_{pq} = \delta_{\sigma_p, \sigma_q} \sum_{\nu} \frac{k_{\nu}^2 \cos(k_{\nu} \cdot r_q - k_{\nu} \cdot r_p)}{2N} \quad U_p = - \sum_{j, \nu \neq 0} \frac{4\pi \zeta_j \cos(k_{\nu} \cdot R_j - k_{\nu} \cdot r_p)}{\Omega k_{\nu}^2} \\ V_{pq} = \sum_{\nu \neq 0} \frac{2\pi \cos(k_{\nu} \cdot r_p - k_{\nu} \cdot r_q)}{\Omega k_{\nu}^2} \end{aligned} \quad (3)$$

where each orbital index p is associated with a spin label $\sigma \in \{\uparrow, \downarrow\}$ and an orbital centroid $r_p = p(2\Omega/N)^{1/d}$ defined in terms of the number of spatial dimensions d and the computational cell volume Ω [28]. The momentum modes are defined as $k_{\nu} = 2\pi\nu/\Omega^{1/d}$ with $\nu \in [-(N/4)^{1/d}, (N/4)^{1/d}]^d$. Throughout, N is the number of spin-orbitals.

We will focus on instances of this Hamiltonian corresponding to periodic materials such as lithium hydride metal, diamond, graphite, crystalline silicon, etc., as well as the special case of $U_p = 0$, which is the uniform electron gas (“jellium”). When dealing with molecular potentials, R_j and ζ_j are the position and charge, respectively, of the j^{th} nucleus. Because these only change Eq. (2) through the terms $U_p Z_p$, they enter our algorithms as a layer of single-qubit rotations per Trotter step, adding essentially no cost to the Trotter step circuit. However, this layer changes the Trotter error and hence the total simulation cost.

The Hubbard model corresponds to the case where T_{pq} are nonzero only for nearest neighbors on a lattice, V_{pq} is the on-site interaction U if p and q correspond to the same orbital (with opposite spin), and $U_p = \mu$ is the chemical potential. Ignoring the chemical potential, the Hubbard Hamiltonian is [37]

$$H = -\tau \sum_{\langle i,j \rangle, \sigma} (a_{i,\sigma}^\dagger a_{j,\sigma} + \text{h.c.}) + U \sum_i n_{i,\uparrow} n_{i,\downarrow}. \quad (4)$$

The first term describes hopping between adjacent lattice sites ($\langle i,j \rangle$ denotes a sum over nearest neighbor lattice sites) and the second is the on-site interaction. Despite its simplicity, the Hubbard model demonstrates a wide range of correlated electron behaviors, including metal-insulator transitions and a superconducting phase, and further is believed to be a candidate model of high-temperature superconductivity in cuprates [38, 39].

Both these Hamiltonians can be simulated using the algorithms of [28] and [29]. The first of these is a split-operator method: it uses a network of Givens rotations or the fast fermionic Fourier transform (FFFT) to alternate between simulating evolution under the kinetic energy terms in the momentum basis and simulating evolution under the potential energy terms in the position basis. (The FFFT was first described in [42] and has also been studied under the name of the “spectral tensor network” [72, 73]; the Givens rotation network was first studied in [43] and further developed in [29]). In the momentum basis, the kinetic energy terms can be simulated using only single-qubit rotations; in the position basis, the potential energy terms can be simulated in linear depth using a swap network [28]. This algorithm operates on a planar array of qubits. By comparison, the second algorithm uses a linear array of qubits, and constructs a swap network which simulates the kinetic and potential operators without the FFFT [29].

We discuss the split-operator algorithm, including the generalization of the fermionic Fourier transform using Givens rotations introduced in [29], in greater depth in Appendix C, and the fermionic swap network in Appendix B. In Appendix C.2 we show how to reduce the cost of simulating the potential energy operator within the split-operator method from $\mathcal{O}(N^2 \log(1/\epsilon))$ to $\mathcal{O}(N^2 + N \log N \log(1/\epsilon))$, where ϵ is the rotation synthesis precision, using Hamming weight phasing. We will discuss in the next subsection the fault-tolerant costs of each Trotter step in terms of the number of arbitrary rotations, Toffoli, and T gates; we describe this in greater detail in Appendix D.

We do not make use of several recent techniques that could lead to improvements in the asymptotic limit, but that have large constant factors and are thus not helpful for the system sizes we consider. For instance, the work of [30] introduced a technique for computing the diagonal part of the plane wave basis electronic structure Hamiltonian with only $\tilde{\mathcal{O}}(N)$ gates; however, the need to compute the discrete Fourier transform of the potential introduces very large constant factors and increases the spatial complexity to $\mathcal{O}(N \log N)$. Haah *et al.* have shown an improved algorithm based on locality of lattice models by dividing simulation into layers of disjoint blocks of evolution (blocks in different layers of the algorithm do overlap), though the number of blocks becomes large in two and especially in three dimensions [45]. Tran *et al.* [74] determined the gate count of the Haah *et al.* algorithm when the interactions decay as a power law, though there is no asymptotic speedup for the Coulomb interaction. Likewise, the work of [46] achieves nearly $N^{o(1)}$ scaling for Hubbard models but by using arbitrarily high-order formulas which increases gate complexity by a factor of 5^k for order k . For the extensive error case we achieve the same scaling using low-order formulas. While randomized Trotter orderings can yield significant improvements for simulation [18, 75], these would come at the cost of the improvements we gain from Hamming weight phasing, which requires specific Trotter orderings for

full utility.

Finally, we note that while the Hamiltonian of Eq. (1) exactly describes any electronic structure system (including molecules) with basis set discretization error asymptotically equivalent to Gaussian-based molecular orbitals, there are other discretization schemes (such as finite difference discretization) which also take the form of Eq. (1), but with different coefficients from those in Eq. (3). Continuing research on basis functions may yield further examples compatible with these algorithms while also being better suited to molecules. For instance, the basis sets described in [76] are compatible with our algorithms while also being significantly more accurate for molecules than plane waves.

2.1 Gate costs per Trotter step

We now summarize how we compute the costs, in terms of the number of arbitrary rotations, Toffoli, and T gates, of the different Trotter steps. We include a more complete description in [Appendix D](#), including upper bounds on the number of the different costly gates, as well as more complete descriptions of the two Trotter steps following [Appendix B](#) and [Appendix C](#). For each Trotter step, we compute these costs numerically. We must calculate this numerically so that we can do it for a range of numbers of ancilla qubits used for Hamming weight phasing. The core idea behind Hamming weight phasing [61] is that the total phasing operation from a group of equiangular rotations can instead be applied to the Hamming weight, using ancilla qubits. The Hamming weight can be computed using only a small number of Toffoli/T gates. We review Hamming weight phasing in greater detail in [Appendix A](#).

Next, we compute the gate costs per Trotter step with fixed numbers of ancilla qubits assigned for Hamming weight phasing. We then determine the numbers of arbitrary rotations that must be synthesized, T, and Toffoli gates, by iterating through the terms in the Hamiltonian in the order specified by the particular Trotter step, combining arbitrary rotations where possible. The number of arbitrary rotations and Toffoli/T gates in each Trotter step are then totally determined by the simulation order, which we describe next. For all systems and for both Trotter steps, we generate the Hamiltonians and determine simulation order using OpenFermion [77].

For the Hubbard model, we simulate all the on-site interactions in a single layer in which we apply Hamming weight phasing. Running through the fermionic swap network (see [Appendix B](#)), we continue until we reverse the initial spin-orbital ordering, applying rotations corresponding to hopping terms as we go. The small number of terms in the Hamiltonian ($O(N^2)$) allows us to defer evolution until we have many equiangular rotations which can be merged using Hamming weight phasing. For jellium (the uniform electron gas), there are too many terms and we must apply these rotations immediately, and can only apply Hamming weight phasing within particular layers of fermionic simulation gates. Finally, for the various materials, we generate the Hamiltonians and run the same simulation procedures as for jellium, but with all the single-qubit rotations repeated in the middle of the Trotter step.

The gate counting for the split-operator Trotter step is slightly easier to describe. For jellium and the materials, we develop a new scheme using the translation-invariance of the interaction part of the Hamiltonian to significantly reduce the Trotter step cost using Hamming weight phasing ([Appendix C.2](#)). For the Hubbard model, because the interaction is exclusively on-site, the split-operator Trotter step simulates it exactly as the fermionic swap network Trotter step does, but without the swaps. In all cases, we group the single-qubit rotations for the kinetic energy operator together as much as possible for Hamming weight phasing given the number of ancilla qubits assigned for that task. Finally, the split-operator step changes from the position to the momentum basis and vice versa: using Hamming weight phasing, arbitrary rotations which might appear in the FFT are catalyzed rather than synthesized (see [Appendix A.3](#)) using only a handful of Toffoli gates, and when the Givens rotation procedure is used we apply Hamming weight phasing within it as much as possible given the number of ancilla qubits.

3 Resource Analysis for Fault-Tolerant Phase Estimation

3.1 Trotter-Suzuki errors

Consistent with past work in molecular simulation [12, 14], we focus on the second-order Trotter formula given by

$$e^{-iHt} \approx \left(\prod_{\ell=1}^L e^{-iH_\ell t/2} \prod_{\ell=L}^1 e^{-iH_\ell t/2} \right) \equiv e^{-iH_{\text{eff}}t}, \quad H = \sum_{\ell=1}^L H_\ell \quad (5)$$

which holds for sufficiently small t . The algorithms of Section 2 naturally implement the second-order Trotter formula if run forward for time $t/2$ and then again in reverse. In Appendix D as well as the previous subsection, we describe how to compute the T-costs of optimized circuits for second-order Trotter steps with both the split-operator and fermionic swap network algorithms. Low-order Trotter decompositions may be useful in variational algorithms to approximate the ground state energy [28, 78]. However, for non-heuristic algorithms using Trotterization to simulate time dynamics or to prepare eigenstates via phase estimation, we must ensure that the discretization errors incurred by the Trotter-Suzuki decomposition can be controlled.

In Appendix E we show that for a second-order Trotter step, the difference between the exact and effective unitary evaluations is bounded by

$$\|e^{-iHt} - e^{-iH_{\text{eff}}t}\| \leq \frac{t^3}{12} \sum_{b=1}^{L-1} \left(\left\| \sum_{c>b} \sum_{a>b} [[H_b, H_c], H_a] \right\| + \frac{1}{2} \left\| \sum_{c>b} [[H_b, H_c], H_b] \right\| \right) \equiv Wt^3 \quad (6)$$

where we call W the ‘‘Trotter error norm’’, and the sums run over the terms in the Hamiltonian. Unlike a similar bound [12, (Eq. 11)], the result above is tighter, and additionally is non-perturbative and holds for all values of t .

Computing the Trotter error norm W allows us to bound the number of applications of the Trotter step circuit, and hence the gate count, required for phase estimation to a desired precision. Of particular interest to us is the problem of sampling eigenvalues using phase estimation. As we show in Appendix E, the maximum shift in the unitary eigenphases (which encode the energies) from Trotterization is

$$|E_n - E_n^{\text{eff}}| t \leq \arctan \left(\|e^{-iHt} - e^{-iH_{\text{eff}}t}\| \frac{\sqrt{4 - \|e^{-iHt} - e^{-iH_{\text{eff}}t}\|^2}}{2 - \|e^{-iHt} - e^{-iH_{\text{eff}}t}\|^2} \right) = Wt^3 + \frac{W^3 t^9}{24} + \mathcal{O}(W^5 t^{15}) \quad (7)$$

for all n , where E_n and E_n^{eff} are corresponding eigenvalues of H and H_{eff} , respectively. Thus, the error in the eigenphases from Trotterization is roughly Wt^3 whenever $Wt^3 \gg W^3 t^9 / 24$. For simplicity, we will assume we can approximate the errors by Wt^3 so long as $Wt^3 \leq 1$, implying the first-order term is at least 24 times larger than the next order correction. As we shall see, the Trotterized evolution time is $t = \sqrt{\Delta E_{TS} / W}$, where ΔE_{TS} accounts for energy errors from the Trotter-Suzuki approximation, so for the condition to hold we need $(\Delta E_{TS})^3 \leq W$.

Rather than approximating the error as Wt^3 , we could instead use the exact expression with the arctan, which holds for $Wt^3 \leq \sqrt{2}$; however, here we focus on the simpler form from the series expansion. The first reason for this is that the condition $Wt^3 \leq 1$ is strongly satisfied for all the numerics in this paper. The second reason is that it is much easier to perform the optimizations of Appendix G if we can assume the error is given by Wt^3 . Unlike the work of [47] which uses Monte Carlo sampling to estimate Trotter errors, we exactly numerically evaluate W for all systems we consider, and are not subject to sampling errors. While the work of [14] performed similar numerics, that work focused on arbitrary basis chemistry (which involved much more complicated error operators), and as a result, was not able to make calculations for larger than $N = 30$, whereas here we are able to go as high as $N = 512$. Still, the bounds in Appendix E can dramatically overestimate the error in the limit

of small t . For this reason, it is important to bear in mind that the estimates we provide for the cost are almost certainly quite pessimistic.

As discussed in [Section 2](#), up to a single layer of rotation gates implementing the local external potential terms, the same circuit simulates Trotter steps of any molecule in the plane wave basis. For simplicity, much of our study focuses on the case when these gates are dropped ($U_p = 0$), corresponding to simulation of the uniform electron gas (jellium) discussed in detail in [\[28\]](#). There, it is argued that the scientific importance of jellium, the classical difficulty of its simulation, and its history as a benchmark for classical electronic structure methods, position it as an intriguing system through which to contrast quantum and classical simulations. There are open questions about the physics of jellium (especially pertaining to the nature of the biased errors introduced to control the sign problem in quantum Monte Carlo) which one could begin to study on a quantum computer with fewer than one hundred logical qubits. Additionally, we compute the Trotter errors after re-introducing the external potential, so as to determine the costs of simulating several different periodic materials.

Jellium has a single phase parameter, the density given by η/Ω in [Eq. \(3\)](#), which scales the system between the limits of strong and weak correlation. Classically, jellium is particularly challenging to study near half-filling for densities where the average electron radius (the Wigner-Seitz radius) is approximately $r_s = 10$ Bohr radii. We conduct most of our numerics for jellium in this regime, computing the Trotter errors for spinful jellium with and without spin in two and three dimensions at $r_s = 10$. Additionally, we compute the Trotter error for jellium in 3D, varying the Wigner-Seitz radius in logarithmically spaced steps from $r_s = 0.5$ to $r_s = 32$ Bohr radii.

For the Hubbard model, we determine the costs of simulation in the intermediate and strongly coupled regimes, $U/\tau = 4$ and $U/\tau = 8$, respectively, with the hopping integral $\tau = 1$. We base our energy precision requirements on classical state-of-the-art ground state energy estimates for the Hubbard model at these parameter values for weakly doped systems (filling fraction 0.875) [\[39\]](#).

For all systems studied, our numerics were performed using code which we have contributed to the open-source package `OpenFermion` [\[77\]](#). However, the larger system sizes studied here may be difficult to access without distributed calculations. We discuss full details of these calculations in [Appendix G](#).

3.2 T gate requirements for Trotterized phase estimation

In order to give concrete T gate counts for the cost of phase estimation, we need to discuss how errors propagate throughout the algorithm. We focus here on quantum simulation within a fault-tolerant architecture, similar to [\[47\]](#), but provide further details about the optimal balance between these errors. Neglecting individual gate errors, the sources of error for estimating the energy in the simulation are:

1. Trotter-Suzuki errors ΔE_{TS} , from the Trotter approximation to e^{-iHt} due to Hamiltonian terms not commuting. From [Eq. \(6\)](#) and [Eq. \(7\)](#), we know that the error in the eigenvalue of the simulated Hamiltonian obeys

$$\Delta E_{TS} \leq Wt^2 + \mathcal{O}(W^3t^8). \quad (8)$$

We numerically computed the Trotter error norm W as shown in [Section 3.1](#) for a variety of systems for the fermionic swap and split-operator algorithms outlined in [Section 2](#).

2. Phase estimation errors ΔE_{PE} , due to uncertainty in the value returned by phase estimation; i.e. errors due to not computing enough bits of the phase. We use the adaptive phase estimation techniques of [\[59, 60\]](#) to reach a root mean squared error of $\Delta E_{PE}t$ using

$$N_{PE} \approx \frac{0.76\pi}{\Delta E_{PE}t} = \frac{0.76\pi\sqrt{W}}{\Delta E_{PE}\sqrt{\Delta E_{TS}}} \quad (9)$$

applications of the fundamental simulation circuit. This relies on using directionally-controlled evolution [\[9\]](#) to reduce prior estimates on the cost of phase estimation [\[59\]](#) by a factor 2. This approach uses a single control qubit. It is possible to match the ultimate lower bound of $\pi/(2\Delta E_{PE}t)$ using multiple control qubits, as in [\[26\]](#). The number of control qubits needed

would be the log of the dynamic range needed. The difference in performance can be accounted for by dividing the gate counts we present by a factor 1.52. Alternatively, the approach of [79] allows median error performance of $\Delta E_{PE}t$ using $1.65/(\Delta E_{PE}t)$ applications of the fundamental simulation circuit and just a single control qubit; this could similarly be accounted for by dividing our gate counts by 1.45.

3. Circuit synthesis errors ΔE_{HT} , due to the fact that there is typically some approximation in compiling arbitrary rotations $R_z(\theta)$ into single-qubit Clifford and T gates. Relative to the other two sources of error, the impact of circuit synthesis error can be reduced at substantially lower cost. If one wishes to synthesize a single-qubit rotation, $R_z(\theta)$ within error ϵ then the number of T gates required using repeat-until-success synthesis is on average $T_{\text{synth}} \approx 1.15 \log_2(1/\epsilon) + 9.2$ [80]. Following the arguments in [47], such errors in the individual rotations add at most linearly to the error in the phase estimated. The phase estimated is the energy eigenvalue multiplied by the time step used for the Trotter decomposition. Thus, if the error desired in the estimate of the energy eigenstate is ΔE_{HT} , the cost of circuit synthesis is approximately

$$N_{HT} \approx 1.15 \log_2 \left(\frac{N_r}{\Delta E_{HT}t} \right) + 9.2 \quad (10)$$

T gates per arbitrary rotation, where N_r is the number of rotations used in a single Trotter step.

In the worst case, these errors add linearly [47]. Thus, to guarantee that the total error is at most ΔE , we assume

$$\Delta E \geq \Delta E_{TS} + \Delta E_{PE} + \Delta E_{HT}. \quad (11)$$

The total T-cost of phase estimation is the product of these, plus the number of “direct” T and Toffoli gates N_d (those T/Toffoli gates which appear in the circuit before synthesis) multiplied by N_{PE} . Again, we convert Toffoli to T gates in computing N_d at a cost of 2 T gates each [81]. Of the different variables described so far, N_r and N_d are the only fixed values: irrespective of the desired total precision ΔE , N_r and N_d are set by the Trotter step circuit and number of ancillae used for Hamming weight phasing. N_{PE} depends on ΔE_{PE} and ΔE_{TS} , and N_{HT} depends directly on ΔE_{HT} and through the time t on ΔE_{TS} . The T count for phase estimation is thus

$$(N_r N_{HT} + N_d) N_{PE} \approx \frac{0.76\pi\sqrt{W}}{\Delta E_{PE}\sqrt{\Delta E_{TS}}} \left(N_r \left[1.15 \log_2 \left(\frac{N_r\sqrt{W}}{\Delta E_{HT}\sqrt{\Delta E_{TS}}} \right) + 9.2 \right] + N_d \right). \quad (12)$$

We numerically minimize this cost for each system, subject to the constraint $0 < \Delta E_{TS} + \Delta E_{PE} + \Delta E_{HT} \leq \Delta E$, where each ΔE is positive. Note that these generally follow the order $\Delta E_{TS} > \Delta E_{PE} \gg \Delta E_{HT}$.

We perform this minimization for two different precision targets: relative precision, where ΔE is set as a fraction of a proxy for the ground state energy \tilde{E}_0 , and thus scales with system size, and absolute precision, where ΔE is set as some fixed constant irrespective of the system size. We discuss these precision targets, as well as some subtleties of the minimization and other numerical results, in more detail in Appendix H. In both cases, in minimizing and plotting the cost, we convert Toffoli to T gates at a cost of 2 T gates per Toffoli. This is because using the techniques of [81], one can distill a Toffoli state at twice the cost of a magic state.

We plot the minimized T-count in Figure 1a for the uniform electron gas in two and three dimensions to relative precision $\Delta E = 0.005 \tilde{E}_0$ (comparable to errors associated with the sign problem in quantum Monte Carlo [82, 83]) and absolute precision $\Delta E = 0.0016$ Hartree (i.e. “chemical accuracy”), which shows that we can realistically expect to simulate classically intractable instances of two- and three-dimensional jellium (up to ~ 300 qubits) with Trotter methods using fewer than one billion T gates. Counts for the Hubbard model to the same relative precision and to absolute precision $\Delta E = \tau/100$ are shown in Figure 2. In Figure 3 we plot the minimized T-counts for phase estimation of several different periodic materials.

We separately plot data for the fermionic swap network and split-operator Trotter steps; within the split-operator Trotter step, we separate data at system side lengths which are a power of two from the rest of the data. When the side length is a power of two, the fermionic fast Fourier transform (FFFT) can be applied allowing significantly more efficient basis changes than the Givens rotation procedure. Because this is not a significant component of the cost for the uniform electron gas, we do not separate these data points for that algorithm. We discuss this difference in costs further in [Appendix C](#). As observed in [\[12, 14, 27, 47\]](#), estimates based on bounds on the Trotter error operator norm are often extremely loose, often by several orders of magnitude [\[14\]](#), so we expect that these values significantly overestimate the true costs. But even if the bounds were tight for the unitary, we are really interested in how the error accumulates for a particular state [\[12, 14\]](#): this difference is yet another potential source of looseness in our cost estimates.

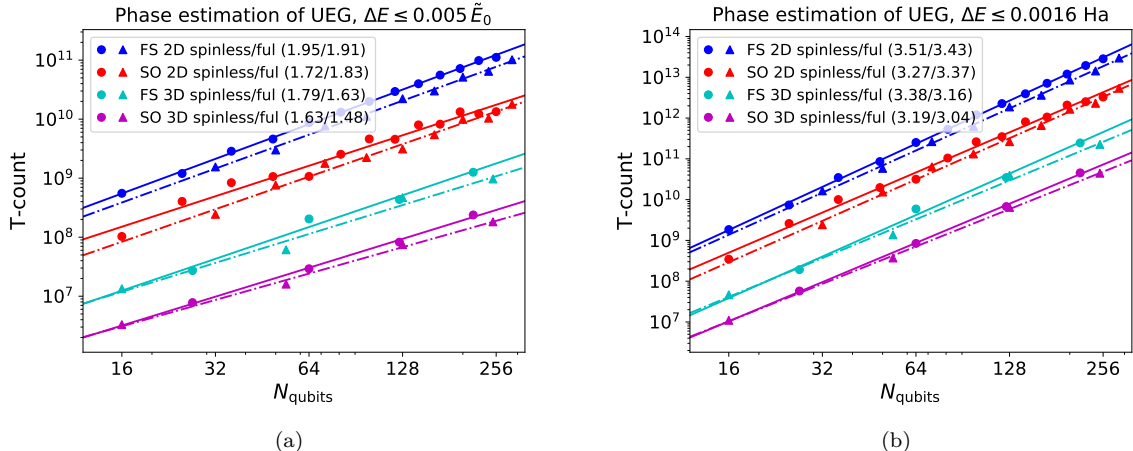


Figure 1: Loose but rigorous upper bound on the number of T gates in the circuit which performs our Trotterized phase estimation simulation of the uniform electron gas (UEG, or jellium) at Wigner-Seitz radius of 10 Bohr radii at half filling $\eta = \lfloor N/2 \rfloor$ based on Eq. (12). (a) shows data for the relative precision target (to within half a percent of the total system energy) as a function of the number of spin-orbitals (qubits); (b) shows data for the absolute precision target of 0.0016 Hartree (“chemical accuracy”). Bounds were computed for jellium both in 2D and in 3D, using both the fermionic swap network (FS) and split-operator (SO) Trotter steps. Exponents of a power law fit are shown in the caption for each system. 14 ancilla qubits are used for Hamming weight phasing for all systems. We can see that for the extensive error target the T cost scales between $\mathcal{O}(N^{3/2})$ and $\mathcal{O}(N^2)$ and for the intensive error target the T cost scales between $\mathcal{O}(N^3)$ and $\mathcal{O}(N^{7/2})$.

An interesting observation from [Figure 2a](#) is that the required number of T gates scales sublinearly in the number of Hamiltonian terms. This is because the number of Trotter steps required to achieve the extensive error target actually decreases with N , implying that the allowable error $\Delta E = \Theta(N)$ grows at a rate comparable to the Trotter error norm. As discussed in [Section 1](#), fixed relative error is often a sensible target. Despite this, for sufficiently large system sizes, the number of circuit repetitions in phase estimation will eventually reach one (it cannot go lower), at which point the number of gates will return to linear scaling in the number of terms in the Hamiltonian. We analyze this further in [Appendix H](#) and show that the number of phase estimation repetitions will not reach this limit until system sizes of hundreds of thousands of spin-orbitals.

The full T-cost (Eq. (12)) scales as $\tilde{\mathcal{O}}(N_r \sqrt{W} / \Delta E^{3/2})$ in terms of the Trotter error norm W , the number of arbitrary rotations per Trotter step N_r , and the target precision ΔE . In [Appendix G](#) we show that the Trotter error norm W scales cubically with N (the number of spin-orbitals) for the uniform electron gas and linearly with N for the Hubbard model. Our relative precision target is $\Delta E = \Theta(N)$ for both the uniform electron gas and the Hubbard model and our absolute precision

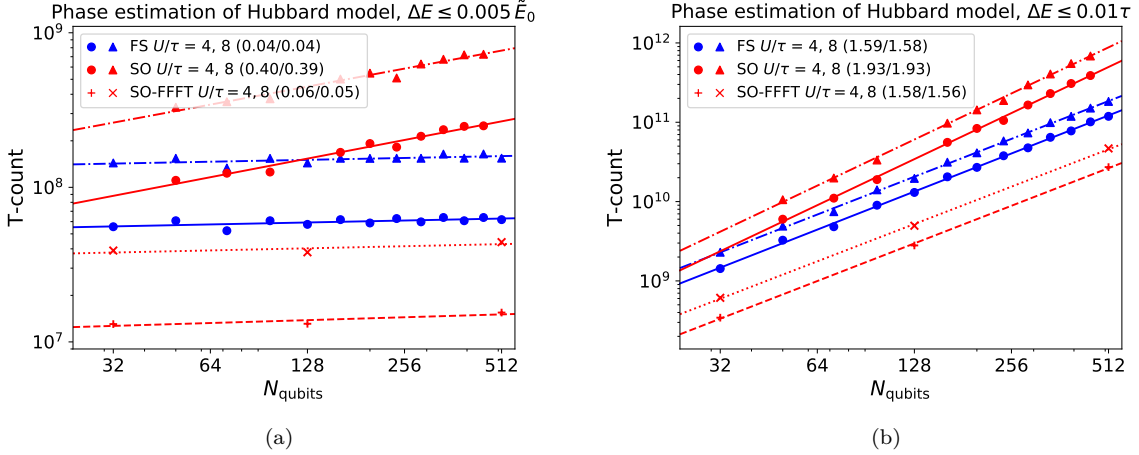


Figure 2: Loose but rigorous upper bound on the number of T gates in the circuit which performs our Trotterized phase estimation simulation of the Hubbard model in the strongly correlated regime ($U/\tau = 4$ and 8) (a) to within half a percent of the total system energy (based on data from [39]), and (b) to absolute precision $\Delta E = \tau/100$. SO-FFFT (red +’s (x’s) for $U/\tau = 4$ (8)) denotes the split-operator Trotter step using the fast fermionic Fourier transform to change bases (only for side lengths which are binary powers), SO (red circles (triangles) for $U/\tau = 4$ (8)) denotes the split-operator Trotter step using Givens rotations to change bases, and FS (blue circles (triangles) for $U/\tau = 4$ (8)) denotes the fermionic swap Trotter step. The fermionic swap Trotter step more efficiently combines repeated rotations for the Hubbard model than for jellium. Lines show power law fits to the data, with the exponents of the fits shown in the caption for each system. 14 ancilla qubits are used for Hamming weight phasing for all systems. We see scaling between $\mathcal{O}(1)$ and $\mathcal{O}(N^{1/2})$ for the extensive error simulations and scaling between $\mathcal{O}(N^{3/2})$ and $\mathcal{O}(N^2)$ for the intensive error simulations.

target is $\Delta E = \Theta(1)$. The number of arbitrary rotations in a Trotter step N_r scales linearly for the Hubbard model with the fermionic swap network Trotter step and either quasilinearly or quadratically for the uniform electron gas and the Hubbard model with the split-operator Trotter step (due to the gates required for changing bases). From these scalings, we expect the T-costs in the absolute precision case to be $\mathcal{O}(N^{7/2})$, $\mathcal{O}(N^2)$, and $\mathcal{O}(N^{3/2})$ for the uniform electron gas, the Hubbard model with fermionic swap network Trotter steps, and the Hubbard model with the split-operator step, respectively. Changing to relative precision reduces these scalings by $N^{3/2}$ in all cases. This is consistent with what we see in Figure 1 and Figure 2.

We can understand the condition from the previous subsection $Wt^3 = \sqrt{(\Delta E_{TS})^3/W} \leq 1$, required for the error to be well estimated by Wt^3 as in Eq. (7), in terms of these scalings. In the absolute precision case the condition becomes more strongly satisfied with growing system size; however, in the relative precision case the quantity Wt^3 will grow. In particular, for the Hubbard model W is only $\mathcal{O}(N)$, so $\sqrt{(\Delta E_{TS})^3/W} = \mathcal{O}(N)$ and we must be careful to work with systems where the value remains small. For the Hubbard model with $U/\tau = 4$, $(\Delta E_{TS})^3/W$ will not reach unity until hundreds of thousands of spin-orbitals, far larger than any system we study.

3.3 Surface code resource estimate analysis

Any resource estimate depends sensitively on the assumed hardware. We shall assume the availability of a large planar square array of qubits with nearest-neighbor interactions capable of implementing all physical gates with non-uniform error rates sufficiently low that the overall performance is approximated by a uniform gate error rate of $p = 10^{-3}$ or $p = 10^{-4}$. We will assume the gates are sufficiently fast to implement a full round of surface code error detection in $1 \mu\text{s}$. Finally, we will assume that a

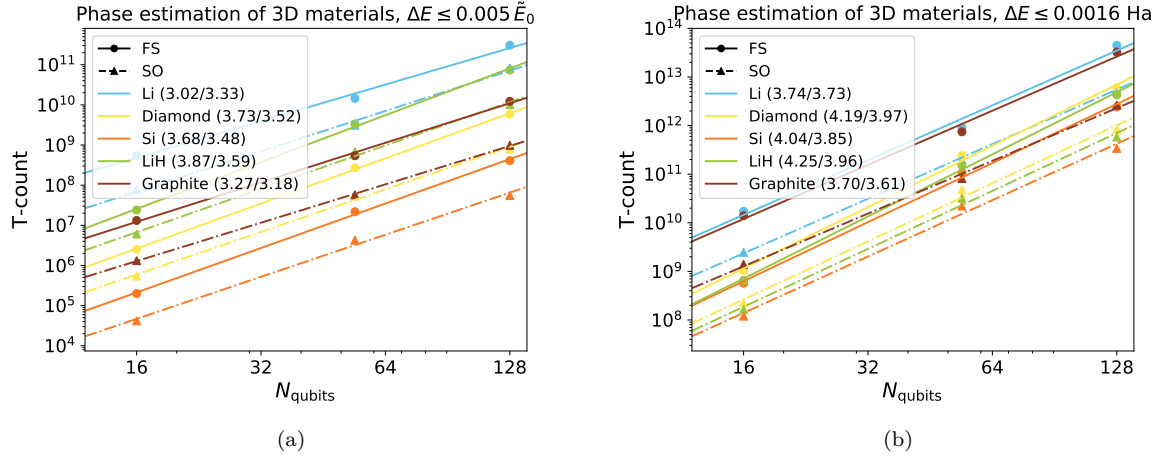


Figure 3: Loose but rigorous upper bound on the number of T gates in the circuit which performs our Trotterized phase estimation simulation for a range of materials (a) to within half a percent of the total system energy, and (b) to absolute precision $\Delta E = 0.0016$ Hartree. Exponents of a power law fit are shown in the legend for each system in the form (x_{FS}/x_{SO}) , where x_{FS} is the exponent for the fermionic swap network Trotter step and x_{SO} is the exponent for the split-operator Trotter step. Lines show power law fits to the data, with the exponents of the fits shown in the caption for each system. 14 ancilla qubits are used for Hamming weight phasing for all systems.

decoder with error suppression performance comparable to [84] is available that is capable of extracting the logical measurement value from potentially many physical measurements associated with a given logical qubit in no more than $10 \mu\text{s}$. This is 100 times slower than has been previously assumed [49], as recent (unpublished) research suggests that it will be challenging to achieve even this level of performance.

We will use lattice surgery [50, 51, 81, 85] to protect our computations, and will make use of distillation to produce $|\text{CCZ}\rangle$ states and catalyzation to efficiently produce $|\text{T}\rangle$ states. These states enable Toffoli and T gates, respectively. Only a single factory producing states serially will be assumed, to minimize qubit overhead. We use the approach described in [81], and fall back to the methods of [51] if a total logical error probability below 0.3 is not attainable during the execution of the phase estimation algorithm. For both techniques, we iterate through level-1 and level-2 distillation code distances of 15 through 51 in steps of 2 to find the scheme and code distances that minimize the total number of physical qubits.

Our resource estimates for phase estimation with the relative precision target (0.5% of the approximate system energy) are in Table 1. We include the total number of logical qubits (including system qubits, and ancilla qubits used for Hamming weight phasing and phase estimation), numbers of Toffoli and T gates, and the computational volume in physical qubit-hours for error rates $p = 10^{-3}$ and $p = 10^{-4}$. We only give results for whichever Trotter step algorithm gives better performance – the fermionic swap network is superior for the Hubbard model when the side length is not a power of two, and the split-operator step superior in the other cases. Table 2 shows the same data with absolute precision targets of 0.0016 Hartree for the uniform electron gas (jellium) and $\tau/100$ for the Hubbard model. The error-corrected data clarify the time-space tradeoff in Hamming weight phasing: by increasing the number of logical ancilla qubits, we can reduce the execution time by as much as an order of magnitude. Because the number of physical qubits used for state distillation is generally comparable to the total number of physical qubits, the increase in the number of physical qubits to support the additional ancillae is typically small, only $\sim 25\%$ of the original requirement.

System	Anc	Log	Toffoli gates	T gates	Physical qubits		Execution time (hours)	
					$p = 10^{-3}$	$p = 10^{-4}$	$p = 10^{-3}$	$p = 10^{-4}$
FH 8×8	4	132	2.5e+05	2.3e+07	3.2e+05 ^b	6.2e+04 ^a	6.0e-01	5.7e-01
FH 8×8	32	160	6.4e+05	9.9e+06	3.3e+05 ^b	7.2e+04 ^a	2.8e-01	3.0e-01
FH 10×10	4	204	1.3e+06	9.2e+07	4.4e+05 ^b	8.8e+04 ^a	2.4e+00	2.3e+00
FH 10×10	128	328	3.1e+06	4.8e+07	6.4e+05 ^b	1.3e+05 ^a	1.4e+00	1.4e+00
FH 16×16	6	518	4.3e+05	2.5e+07	9.4e+05 ^b	2.0e+05 ^a	6.6e-01	6.4e-01
FH 16×16	130	642	9.0e+05	9.2e+06	1.1e+06 ^b	2.5e+05 ^a	2.8e-01	3.1e-01
UEG $3 \times 3 \times 3$	4	58	4.4e+05	3.1e+07	1.9e+05 ^b	3.5e+04 ^a	8.2e-01	7.8e-01
UEG $3 \times 3 \times 3$	16	70	1.0e+06	1.4e+07	2.0e+05 ^b	4.0e+04 ^a	4.0e-01	4.3e-01
UEG $4 \times 4 \times 4$	4	132	2.0e+06	1.5e+08	3.3e+05 ^b	6.2e+04 ^a	4.0e+00	3.8e+00
UEG $4 \times 4 \times 4$	128	256	5.7e+06	2.5e+07	5.2e+05 ^b	1.1e+05 ^a	8.9e-01	1.1e+00
UEG $5 \times 5 \times 5$	4	254	5.0e+06	4.0e+08	6.0e+05 ^b	1.4e+05 ^a	1.1e+01	1.0e+01
UEG $5 \times 5 \times 5$	128	378	1.4e+07	3.3e+07	7.3e+05 ^b	2.1e+05 ^a	1.5e+00	2.1e+00

Table 1: Fault-tolerant resource estimates (number of logical ancillae (Anc), logical qubits for encoding the system (Log), Toffoli gates, T gates, physical qubits, and execution time) for the relative precision target, assuming physical operation error rates of $p = 10^{-3}$ and $p = 10^{-4}$. We include data varying the number of ancilla qubits used for Hamming weight phasing. FH indicates Hubbard model data with $U/\tau = 4$, and UEG indicates data for the spinful uniform electron gas with a Wigner-Seitz radius of 10 Bohr radii. For all systems, the total number of logical qubits for the system (Log) is double the number of grid points plus the number of ancillae (Anc). When generating T states, superscript *a* indicates that a single round of 15-1 distillation was used, superscript *b* indicates 15-1 followed by CCZ state generation and catalyzation was used [81]. Scheme *a* requires fewer qubits, but can be slower than scheme *b*.

4 Discussion

Prior to this work, the most practical techniques for performing fault-tolerant quantum simulations of condensed-phase correlated electron models were those introduced in [26]. That work deployed methods based on linear combinations of unitaries which scaled roughly as $\mathcal{O}(\lambda N/\Delta E)$ where λ is the sum of the absolute values of the Hamiltonian terms; $\lambda = \mathcal{O}(N)$ for the Hubbard model and $\lambda = \mathcal{O}(N^2)$ for the plane wave basis electronic structure Hamiltonian. In this paper we demonstrate low-order Trotter methods for these problems scaling as roughly $\mathcal{O}(N^{3/2}/\Delta E^{3/2})$ for the Hubbard model and $\mathcal{O}(N^{7/2}/\Delta E^{3/2})$ for jellium. While the methods of [26] are more practical when targeting an intensive error, $\Delta E = \Theta(1)$, the methods here are more practical when targeting an extensive error, $\Delta E = \Theta(N)$. For example, for Hubbard model simulations with extensive error the methods of [26] have $\mathcal{O}(N)$ T complexity whereas the methods here have $\mathcal{O}(1)$ T complexity (at least until $N > 10^5$ at which point the scaling increases).

It is also interesting to compare these asymptotics to other recent work on Trotter based simulations. For instance, the work of [46] will also have effective $\mathcal{O}(1)$ scaling for the Hubbard model in the extensive error target regime, but does not provide a concrete circuit implementation and their analysis is not directly applicable to electronic structure Hamiltonians. Another comparison is the method of [18] which essentially scales as $\mathcal{O}(\lambda^2/\Delta E^2)$ for the purposes of phase estimation. For extensive error simulations that method will have T complexities of roughly $\mathcal{O}(N^2)$ and $\mathcal{O}(1)$ for jellium and the Hubbard model, respectively, which matches our scaling here.

Our work has also made significant progress towards reducing the resources required to perform the first classically intractable electronic structure calculation on a fault-tolerant device that could be realized with the error rates and connectivities accessible to superconducting qubits. In addition to the low scaling of our methods, as discussed above, we put significant effort into optimizing constant factors through strategies such as the Hamming weight phasing techniques introduced here. We also further reduced surface code compilation overheads by using lattice surgery [51, 81, 85] rather than topological braiding (as in [26]). Ultimately, we found that even at 10^{-3} error rates one can error-

System	Anc	Log	Toffoli gates	T gates	Physical qubits		Execution time (hours)	
					$p = 10^{-3}$	$p = 10^{-4}$	$p = 10^{-3}$	$p = 10^{-4}$
FH 8×8	4	132	4.6e+07	5.0e+09	4.2e+05 ^b	8.1e+04 ^a	1.3e+02	1.2e+02
FH 8×8	32	160	1.2e+08	2.1e+09	4.7e+05 ^b	9.6e+04 ^a	5.9e+01	6.2e+01
FH 10×10	4	204	4.7e+08	4.2e+10	6.7e+05 ^b	2.4e+05 ^b	1.2e+03	1.1e+03
FH 10×10	128	328	1.1e+09	2.2e+10	9.8e+05 ^b	3.2e+05 ^b	7.0e+02	6.2e+02
FH 16×16	6	518	6.4e+08	4.7e+10	1.7e+06 ^b	4.5e+05 ^b	1.4e+03	1.2e+03
FH 16×16	130	642	1.3e+09	1.6e+10	1.8e+06 ^b	5.3e+05 ^b	5.3e+02	4.7e+02
UEG $3 \times 3 \times 3$	4	58	8.9e+06	7.2e+08	2.3e+05 ^b	4.4e+04 ^a	1.9e+01	1.8e+01
UEG $3 \times 3 \times 3$	16	70	2.0e+07	3.2e+08	2.5e+05 ^b	5.0e+04 ^a	9.1e+00	9.6e+00
UEG $4 \times 4 \times 4$	4	132	1.4e+08	1.3e+10	4.8e+05 ^b	1.7e+05 ^b	3.9e+02	3.4e+02
UEG $4 \times 4 \times 4$	128	256	4.2e+08	2.2e+09	6.9e+05 ^b	1.4e+05 ^a	7.5e+01	9.3e+01
UEG $5 \times 5 \times 5$	4	254	1.0e+09	9.9e+10	8.9e+05 ^c	2.7e+05 ^b	5.1e+03	2.6e+03
UEG $5 \times 5 \times 5$	128	378	2.8e+09	8.4e+09	1.1e+06 ^b	3.6e+05 ^b	3.8e+02	3.4e+02

Table 2: Fault-tolerant resource estimates (number of logical ancillae (Anc), logical qubits for encoding the system (Log), Toffoli gates, T gates, physical qubits, and execution time) for the absolute precision target, assuming physical operation error rates of $p = 10^{-3}$ and $p = 10^{-4}$. We include data varying the number of ancilla qubits used for Hamming weight phasing. FH indicates Hubbard model data with $U/\tau = 4$, and UEG indicates data for the spinful uniform electron gas with a Wigner-Seitz radius of 10 Bohr radii. For all systems, the total number of logical qubits for the system (Log) is double the number of grid points in it plus the number of ancillae (Anc). When generating T states, superscript *a* indicates a single round of 15-1 distillation was used, superscript *b* indicates 15-1 followed by CCZ state generation and catalyzation was used, superscript *c* indicates two rounds of 15-1 distillation were used [81]. Note that scheme *a*, while requiring fewer qubits, can be slower than scheme *b*.

correct classically intractable instances of these problems with only a few hundred thousand physical qubits.

We conclude with a brief discussion of open research questions relevant to this work. First, one could investigate the T-costs of higher-order Trotter formulas. Other work has shown these to be advantageous even for small system sizes [52, 53]. Numerically evaluating the Trotter error norm would be challenging for higher-order formulas, but even analytic bounds would be interesting to compare with those presented here. Second, rather than looking at upper bounds, one could study the true Trotter error in circuit simulations (using inefficient exact classical calculations and extrapolating, or with efficient approximate classical calculations): this could reduce the number of Trotter steps, and hence our T-count estimates, by orders of magnitude. Finally, one could attempt to extrapolate these quantum simulations to zero Trotter error with even fewer steps than we use here by interpolating the trend between Trotter errors and Trotter number, as is common in quantum Monte Carlo electronic structure simulations.

Acknowledgments

We thank Garnet Kin-Lic Chan, Matthew Foulkes, Jessica Lemieux, Matthias Degroote, and Kostyantyn Kechedzhi for discussions about when it is appropriate to target extensive error in the energies of Hubbard model and jellium simulations, and for what values of extensive and intensive errors the simulations would be classically challenging. We thank Vasil Denchev for assistance parallelizing some calculations required for our numerics, as well as contributors to OpenFermion [77], which we used extensively for numerics. We thank Yuan Su for feedback on an earlier draft. I. D. K. acknowledges partial support from the National Sciences and Engineering Research Council of Canada. D. W. B. is funded by Australian Research Council Discovery projects (Grant Nos. DP160102426 and DP190102633) and by a grant from Google Quantum.

References

- [1] R. P. Feynman, *International Journal of Theoretical Physics* **21**, 467 (1982).
- [2] S. Lloyd, *Science* **273**, 1073 (1996).
- [3] D. S. Abrams and S. Lloyd, *Physical Review Letters* **83**, 5162 (1999).
- [4] A. Y. Kitaev, *arXiv preprint arXiv:9511026* (1995).
- [5] D. S. Abrams and S. Lloyd, *Physical Review Letters* **79**, 2586 (1997).
- [6] N. M. Tubman, C. Mejuto-Zaera, J. M. Epstein, D. Hait, D. S. Levine, W. Huggins, Z. Jiang, J. R. McClean, R. Babbush, M. Head-Gordon, and K. B. Whaley, *arXiv preprint arXiv:1809.05523* (2018).
- [7] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).
- [8] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, *Molecular Physics* **109**, 735 (2011).
- [9] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, *Phys. Rev. A* **90**, 022305 (2014).
- [10] H. F. Trotter, *Proceedings of the American Mathematical Society* **10**, 545 (1959).
- [11] M. Suzuki, *Journal of Mathematical Physics* **32**, 400 (1991).
- [12] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, *Quantum Info. Comput.* **15**, 361 (2015).
- [13] J. R. McClean, R. Babbush, P. J. Love, and A. Aspuru-Guzik, *The Journal of Physical Chemistry Letters* **5**, 4368 (2014).
- [14] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, *Phys. Rev. A* **91**, 022311 (2015).
- [15] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, *Quantum Info. Comput.* **15**, 1 (2015).
- [16] F. Motzoi, M. Kaicher, and F. Wilhelm, *Physical review letters* **119**, 160503 (2017).
- [17] M. Motta, E. Ye, J. R. McClean, Z. Li, A. J. Minnich, R. Babbush, and G. K.-L. Chan, *arXiv preprint arXiv:1808.02625* (2018).
- [18] E. Campbell, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [19] A. M. Childs and N. Wiebe, *Quantum Info. Comput.* **12**, 901 (2012).
- [20] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [21] R. Babbush, D. W. Berry, I. D. Kivlichan, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, *New Journal of Physics* **18**, 033032 (2016).
- [22] G. H. Low and I. L. Chuang, *Quantum* **3**, 163 (2019).
- [23] G. H. Low and I. L. Chuang, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [24] D. Poulin, A. Kitaev, D. S. Steiger, M. B. Hastings, and M. Troyer, *Phys. Rev. Lett.* **121**, 010501 (2018).
- [25] D. W. Berry, M. Kieferová, A. Scherer, Y. R. Sanders, G. H. Low, N. Wiebe, C. Gidney, and R. Babbush, *npj Quantum Information* **4**, 22 (2018).
- [26] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, *Physical Review X* **8**, 041015 (2018).
- [27] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, *Quantum* **3**, 208 (2019).
- [28] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven, and G. K.-L. Chan, *Phys. Rev. X* **8**, 011044 (2018).
- [29] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [30] G. H. Low and N. Wiebe, *arXiv preprint arXiv:1805.00675* (2018).
- [31] C. Zalka, *Fortschritte der Physik* **46**, 877 (1998).
- [32] B. Toloui and P. J. Love, *arXiv preprint arXiv:1312.2579* (2013).
- [33] R. Babbush, D. W. Berry, Y. R. Sanders, I. D. Kivlichan, A. Scherer, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, *Quantum Science and Technology* **3**, 015006 (2018).
- [34] I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik, *Proceedings of the National Academy of Sciences* **105**, 18681 (2008).

- [35] I. D. Kivlichan, N. Wiebe, R. Babbush, and A. Aspuru-Guzik, *Journal of Physics A: Mathematical and Theoretical* **50**, 305301 (2017).
- [36] R. Babbush, D. W. Berry, J. R. McClean, and H. Neven, *npj Quantum Information* **5**, 92 (2019).
- [37] J. Hubbard, *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **276**, 238 (1963).
- [38] P. A. Lee, N. Nagaosa, and X.-G. Wen, *Rev. Mod. Phys.* **78**, 17 (2006).
- [39] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu, and E. Gull (Simons Collaboration on the Many-Electron Problem), *Phys. Rev. X* **5**, 041041 (2015).
- [40] G. Ortiz, J. Gubernatis, E. Knill, and R. Laflamme, *Physical Review A* **64**, 022319 (2001).
- [41] F. Verstraete and J. I. Cirac, *Journal of Statistical Mechanics: Theory and Experiment* **2005**, P09012 (2005).
- [42] F. Verstraete, J. I. Cirac, and J. I. Latorre, *Phys. Rev. A* **79**, 032316 (2009).
- [43] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, *Physical Review A* **92**, 062318 (2015).
- [44] Z. Jiang, K. J. Sung, K. Kechedzhi, V. N. Smelyanskiy, and S. Boixo, *Physical Review Applied* **9**, 044036 (2018).
- [45] J. Haah, M. Hastings, R. Kothari, and G. H. Low, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2018) pp. 350–360.
- [46] A. M. Childs and Y. Su, *Phys. Rev. Lett.* **123**, 050503 (2019).
- [47] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, *Proceedings of the National Academy of Sciences* **114**, 7555 (2017).
- [48] R. Babbush, D. W. Berry, and H. Neven, *Physical Review A* **99**, 040301 (2019).
- [49] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Physical Review A* **86**, 32324 (2012).
- [50] C. Horsman, A. G. Fowler, S. Devitt, and R. V. Meter, *New Journal of Physics* **14**, 123011 (2012).
- [51] A. G. Fowler and C. Gidney, *arXiv preprint arXiv:1808.06709* (2018).
- [52] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, *Proceedings of the National Academy of Sciences* **115**, 9456 (2018).
- [53] Y. Nam and D. Maslov, *npj Quantum Information* **5**, 44 (2019).
- [54] P. Corboz, *Phys. Rev. B* **93**, 045116 (2016).
- [55] J. J. Shepherd and A. Grüneis, *Phys. Rev. Lett.* **110**, 226401 (2013).
- [56] J. J. Shepherd, T. M. Henderson, and G. E. Scuseria, *Phys. Rev. Lett.* **112**, 133002 (2014).
- [57] P.-F. Loos and P. M. W. Gill, *Wiley Interdisciplinary Reviews: Computational Molecular Science* **6**, 410 (2016).
- [58] J. McClain, J. Lischner, T. Watson, D. A. Matthews, E. Ronca, S. G. Louie, T. C. Berkelbach, and G. K.-L. Chan, *Phys. Rev. B* **93**, 235139 (2016).
- [59] D. W. Berry, B. L. Higgins, S. D. Bartlett, M. W. Mitchell, G. J. Pryde, and H. M. Wiseman, *Phys. Rev. A* **80**, 052114 (2009).
- [60] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde, *Nature* **450**, 393 (2007).
- [61] C. Gidney, *Quantum* **2**, 74 (2018).
- [62] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, *Phys. Rev.* **108**, 1175 (1957).
- [63] P. Jordan and E. Wigner, *Zeitschrift für Physik* **47**, 631 (1928).
- [64] R. D. Somma, G. Ortiz, J. Gubernatis, E. Knill, and R. Laflamme, *Phys. Rev. A* **65**, 17 (2002).
- [65] S. Bravyi and A. Kitaev, *Annals of Physics* **298**, 210 (2002).
- [66] J. T. Seeley, M. J. Richard, and P. J. Love, *Journal of Chemical Physics* **137**, 224109 (2012).
- [67] A. Tranter, S. Sofia, J. Seeley, M. Kaicher, J. McClean, R. Babbush, P. V. Coveney, F. Mintert, F. Wilhelm, and P. J. Love, *International Journal of Quantum Chemistry* **115**, 1431 (2015).

- [68] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, [arXiv preprint arXiv:1701.08213 \(2017\)](#).
- [69] J. D. Whitfield, V. Havlíček, and M. Troyer, [Physical Review A **94**, 030301 \(2016\)](#).
- [70] V. Havlíček, M. Troyer, and J. D. Whitfield, [Phys. Rev. A **95**, 032332 \(2017\)](#).
- [71] Z. Jiang, J. McClean, R. Babbush, and H. Neven, [Phys. Rev. Applied **12**, 064041 \(2019\)](#).
- [72] A. J. Ferris, [Phys. Rev. Lett. **113**, 010401 \(2014\)](#).
- [73] A. Chandran, J. Carrasquilla, I. H. Kim, D. A. Abanin, and G. Vidal, [Physical Review B **92**, 024201 \(2015\)](#).
- [74] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, [Phys. Rev. X **9**, 031006 \(2019\)](#).
- [75] A. M. Childs, A. Ostrander, and Y. Su, [Quantum **3**, 182 \(2019\)](#).
- [76] S. R. White, [The Journal of Chemical Physics **147**, 244102 \(2017\)](#).
- [77] J. McClean, N. Rubin, K. Sung, I. D. Kivlichan, X. Bonet-Monroig, Y. Cao, C. Dai, E. S. Fried, C. Gidney, B. Gimby, P. Gokhale, T. Haner, T. Hardikar, V. Havlek, O. Higgott, C. Huang, J. Izaac, Z. Jiang, X. Liu, S. McArdle, M. Neeley, T. O’Brien, B. O’Gorman, I. Ozfidan, M. D. Radin, J. Romero, N. P. D. Sawaya, B. Senjean, K. Setia, S. Sim, D. S. Steiger, M. Steudtner, Q. Sun, W. Sun, D. Wang, F. Zhang, and R. Babbush, [Quantum Science and Technology \(2020\)](#), [10.1088/2058-9565/ab8ebc](#).
- [78] D. Wecker, M. B. Hastings, and M. Troyer, [Phys. Rev. A **92**, 042303 \(2015\)](#).
- [79] N. Wiebe and C. Granade, [Phys. Rev. Lett. **117**, 010503 \(2016\)](#).
- [80] A. Bocharov, M. Roetteler, and K. M. Svore, [Phys. Rev. Lett. **114**, 080502 \(2015\)](#).
- [81] C. Gidney and A. G. Fowler, [Quantum **3**, 135 \(2019\)](#).
- [82] B. Tanatar and D. M. Ceperley, [Phys. Rev. B **39**, 5005 \(1989\)](#).
- [83] J. J. Shepherd, G. Booth, A. Grüneis, and A. Alavi, [Phys. Rev. B **85**, 081103 \(2012\)](#).
- [84] A. G. Fowler, [arXiv preprint arXiv:1310.0863 \(2013\)](#).
- [85] D. Litinski, [Quantum **3**, 128 \(2019\)](#).
- [86] J. E. Savage, *Models of computation*, Vol. 136 (Addison-Wesley Reading, MA, 1998).
- [87] J. Nordstrom, [Logical Methods in Computer Science **Volume 9**, **Issue 3** \(2013\)](#), [10.2168/LMCS-9\(3:15\)2013](#).
- [88] R. Bhatia and C. Davis, [Linear and Multilinear Algebra **15**, 71 \(1984\)](#).
- [89] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, *et al.*, [Wiley Interdisciplinary Reviews: Computational Molecular Science **8**, e1340 \(2018\)](#).
- [90] T. Chachiyo, [The Journal of Chemical Physics **145**, 021101 \(2016\)](#).

A Combining repeated rotations using Hamming weight phasing

In this appendix, we discuss the technique of [61] for efficiently applying repeated arbitrary rotations by the same angle. First, we review how this may be done when the repeated rotations by the same angle are parallelizable, that is, multiple rotations $R_z(\theta)$ can be performed simultaneously in the circuit. For n parallel equiangular rotations, the technique reduces the number of arbitrary rotations which must be synthesized to $\lceil \log_2 n + 1 \rceil$, at the cost of at most $4n - 4$ additional T gates (or $n - 1$ Toffoli gates) and $n - 1$ ancilla qubits. Because synthesizing arbitrary rotations is very costly on a fault-tolerant quantum computer, this can lead to large savings. Following that, we briefly discuss how the method works with a limited number of ancillae $r < n - 1$ as well as some possible extensions.

A.1 Combining arbitrary parallelizable rotations by Hamming weight phasing

Consider a circuit where the same single-qubit rotation $R_z(\theta) = \exp(-i\theta Z/2)$ is simultaneously applied to three different qubits. The action on the logical states is a different phase, depending only on the Hamming weight of that logical state. To be explicit,

1. the all-zero state $|000\rangle$ picks up the phase $-3\theta/2$,
2. the three states of Hamming weight one $|001\rangle$, $|010\rangle$, and $|100\rangle$ each pick up a phase $-\theta/2$
3. the three states of Hamming weight two $|011\rangle$, $|101\rangle$, and $|110\rangle$ are all phased by $\theta/2$, and finally
4. the all-one state $|111\rangle$ picks up the phase $3\theta/2$.

Rather than applying three rotations by the same angle $R_z(\theta)$, we can instead compute the Hamming weight of the input states, and apply two distinct rotations to the Hamming weight: $R_z(\theta)$ to the 1s bit, and $R_z(2\theta)$ to the 2s bit. In this case,

1. the phase $-3\theta/2$ is applied to the all-zero state $|000\rangle$ (Hamming weight 0),
2. $-\theta/2 = -\theta + \theta/2$ is applied to $|001\rangle$, $|010\rangle$, and $|100\rangle$ (Hamming weight 1 = 01_b in binary),
3. $\theta/2 = +\theta - \theta/2$ is applied to $|011\rangle$, $|101\rangle$, and $|110\rangle$ (Hamming weight 2 = 10_b), and finally
4. $3\theta/2 = +\theta + \theta/2$ is applied to the all-one state $|111\rangle$ (Hamming weight 3 = 11_b).

The phases on each logical state are identical for the two procedures. However, because arbitrary rotations must be synthesized using costly T gates, reducing the number of rotation gates in the circuit reduces its fault-tolerant cost.

This idea readily generalizes to the case of n repeated equiangular rotations $R_z(\theta)$ appearing in parallel in a circuit: rather than applying the n original arbitrary rotations, we can compute the Hamming weight of the relevant qubits, and instead apply $\lfloor \log_2 n + 1 \rfloor$ arbitrary rotations $R_z(\theta)$, $R_z(2\theta)$, $R_z(4\theta)$, \dots , to the Hamming weight [61]. We call this technique Hamming weight phasing.

At this point, Hamming weight phasing would appear to give an improvement for free. However, there are two sources of cost we have not mentioned. These costs are an additional number of T gates and ancilla qubit requirement, both arising from the use of adder circuits to compute the Hamming weight. Using the adder circuits of [61], we can compute and uncompute the Hamming weight using at most $n - 1$ Toffoli gates (or equivalently, $4n - 4$ additional T gates) and $n - 1$ ancilla qubits, corresponding to $n - 1$ single-bit adders. (Each single-bit adder uses one ancilla qubit, and costs one Toffoli gate or 4 T gates [61].) The strategy for this summation is as follows.

We begin with n qubits to which a rotation $R_z(\theta)$ must be applied. We call these initial qubits “weight-1” qubits: our goal is to combine them such that we are left with one qubit which captures each digit of the Hamming weight, i.e., one qubit each of weight 1, 2, 4, 8, \dots , $2^{\lfloor \log_2(n) \rfloor}$. Then, we apply the rotation $R_z(w\theta)$ to each qubit depending on its weight w , and uncompute the Hamming weight, returning each qubit to its initial state.

We combine the qubits to form the Hamming weight in stages. First, we pick a group of 3 weight-1 qubits: call these three qubits a , b , and c . We initialize one ancilla qubit in the $|0\rangle$ state, and apply the adder building block of [61] to the three qubits and the ancilla. The first two qubits are unaffected (their final state remains a and b), but the third qubit and the ancilla qubit are changed to the 1s and 2s bits of $(a + b + c)$: $(a + b + c)_0$ and $(a + b + c)_1$, respectively. After the adder, we have introduced one ancilla qubit (now a weight-2 qubit, storing $(a + b + c)_1$), and accounted for the weight-1 qubits a and b , thereby decreasing the number of weight-1 qubits unaccounted for in the Hamming weight by two. We repeat this process of choosing groups of 3 weight-1 qubits and combining them until either two or one weight-1 qubits remain. At least one weight-1 qubit must remain because the adder primitive always leaves a qubit of weight 1.

At this point, we have applied $\lfloor (n-1)/2 \rfloor$ temporary adders, thereby introducing $\lfloor (n-1)/2 \rfloor$ ancilla qubits (now of weight 2) and incurring a cost of $\lfloor (n-1)/2 \rfloor$ Toffoli gates. $n - 2\lfloor (n-1)/2 \rfloor$ weight-1 qubits remain. If n is odd, only one weight-1 qubit remains, and we are done: this is the 1s digit qubit in the Hamming weight. If n is even, two weight-1 qubits (call them a and b) still have to be accounted for. We introduce an ancilla and apply the temporary adder to the two remaining weight-1 qubits. We are left with the input a , $(a + b)_0$ —now the weight-1 bit of the Hamming weight—and the

ancilla in the state $(a + b)_1$, which now has weight 2. In the worst case (the case where n is even), we applied $\lfloor n/2 \rfloor$ adders, introduced $\lfloor n/2 \rfloor$ ancilla qubits, and used $\lfloor n/2 \rfloor$ Toffoli gates to reduce to a single weight-1 qubit. We repeat this procedure for the weight-2 qubits. There are at most $\lfloor n/2 \rfloor$ weight-2 qubits, so the cost of reducing to a single weight-2 qubit for the Hamming weight register is at most $n/4$ temporary adders.

In general, when converting from weight w to weight $2w$, the odd case is simple. In the even case, two weight- w qubits (again, a and b) remain, which are converted to a , $(a + b)_{\log w}$, and $(a + b)_{\log w + 1}$; these are the unmodified input, the weight- w bit of the Hamming weight, and the ancilla qubit now in the weight $2w$ state $(a + b)_{\log w + 1}$, respectively. We continue to repeat this procedure for the higher-weight qubits, determining the Hamming weight bit-by-bit, until we reach a stage where there are exactly two qubits of weight $2^{\lfloor \log_2 n \rfloor - 1}$. At this point, we apply the final adder, and introduce a final ancilla qubit for the most significant bit of the Hamming weight. The worst case for the entire procedure is when n is a power of 2. When n is a power of 2, the number of weight- k qubits is always even up to the final weight to be combined, weight $k = 2^{\lfloor \log_2 n \rfloor - 1} = n/2$, for which there are 2 qubits. In general, the number of qubits of weight k is n/k for each k , except for the single $k = n$ qubit. So in the worst case the number of adders required is half this expression summed over k ,

$$\sum_{i=0}^{\lfloor \log_2 n \rfloor - 1} \frac{n}{2^{i+1}} = n - 1. \quad (13)$$

The total cost to form the Hamming weight register is thus at most $n - 1$ Toffoli gates (equivalently, $4n - 4$ T gates) and $n - 1$ ancilla. This allows us to reduce the number of parallel equiangular rotations from n to $\lfloor \log_2 n + 1 \rfloor$. Uncomputing the Hamming weight requires no further T gates or ancillae [61].

Finally, note if $\theta = m\pi/2^k$ for integers k and m that there is no reason to continue the addition past weight 2^{k-2} : at that point, each qubit only needs a T gate applied to it, and further adders eliminate only one T gate, versus the 4 T gates used by the adder. We discuss a specific application of this idea later in [Appendix A.3](#).

The total tradeoff for Hamming weight phasing depends on the cost of synthesizing arbitrary rotations. With T_{synth} the number of T gates required to synthesize a single-qubit arbitrary rotation $R_z(\theta)$, Hamming weight phasing reduces the number of T gates required by

$$T_{\text{synth}}(n - \lfloor \log_2 n + 1 \rfloor) - 4(n - 1). \quad (14)$$

Even for small n , the break-even point for Hamming weight phasing to yield an improvement is $T_{\text{synth}} \gtrsim 10$. For comparison, the number of T gates required to synthesize the single-qubit rotation $R_z(\theta)$ within error ϵ_{synth} using repeat-until-success circuits is approximately $T_{\text{synth}} \approx 1.15 \log_2(1/\epsilon_{\text{synth}}) + 9.2$ [80].

The additional cost that must be considered is that of allocating $n - 1$ ancilla qubits for use in Hamming weight phasing. If n is sufficiently large, these qubits could alternatively be put to use distilling T gates. Together with the additional space requirement, this must be balanced against the savings of Hamming weight phasing.

In the following subsection, we consider the problem of Hamming weight phasing with stronger restrictions on the number of ancilla qubits we can introduce. We present two schemes: one with a constant number of ancillae r , and one using $\lfloor \sqrt{n} \rfloor + \lfloor \log_2 n \rfloor$ ancilla qubits. In both cases, the T gate requirement increases while the ancilla qubit requirement becomes more favourable.

A.2 Hamming weight phasing with limited ancilla

The ancilla requirement of Hamming weight phasing may be prohibitive. We discuss two possible modifications with lower ancilla requirements in this subsection. First, we consider Hamming weight phasing when limiting to a constant number of ancilla qubits $r < n - 1$. Second, we discuss a method where we compute the Hamming weight of size- $\lfloor \sqrt{n} \rfloor$ subsets of the qubits at a time, and then sum those subset Hamming weights, rather than directly computing the full Hamming weight. This second

method reduces the number of ancilla qubits required to $\lfloor \sqrt{n} \rfloor + \lfloor \log_2 n \rfloor$. For both modifications, the number of ancilla is reduced at the cost of requiring more Toffoli or T gates.

First, we consider limiting to a constant number of ancilla qubits $r < n - 1$. Because the number of parallel equiangular rotations n is greater than $r + 1$, it is no longer possible to combine all the rotations at once. The problem is that we do not have enough ancilla qubits to compute the Hamming weight. However, we can still combine $r + 1$ rotations at a time into $\lfloor \log_2(r + 1) \rfloor + 1$ arbitrary rotations, and repeat the process until we have combined all n arbitrary rotations. Within each repetition, there are at most r adders using all r ancilla qubits, reducing $r + 1$ arbitrary rotations to $\lfloor \log_2(r + 1) \rfloor + 1$ at a cost of r Toffoli ($4r$ T) gates. The number of repetitions required to include every rotation is $\lceil n/(r + 1) \rceil$. The total cost is at most $4r \lceil n/(r + 1) \rceil$ T gates and $\lceil n/(r + 1) \rceil \lfloor \log_2(r + 1) \rfloor + 1$ arbitrary rotations. So even with only a constant r ancilla qubits, we reduce the number of T gates by at least

$$T_{\text{synth}} \left(n - \left\lceil \frac{n}{r + 1} \right\rceil \lfloor \log_2(r + 1) \rfloor + 1 \right) - 4r \left\lceil \frac{n}{r + 1} \right\rceil, \quad (15)$$

where T_{synth} is the number of T gates per arbitrary rotation required for synthesis. This is the simplest approach to reducing the number of rotations with a limited number of ancillae.

Second, we present a method for reducing the number of arbitrary rotations by computing the Hamming weight of groups of qubits, and adding it to a second register storing the total Hamming weight. We divide the original n rotations into groups of size $\lfloor \sqrt{n} \rfloor$. For each group, we compute the Hamming weight using $\lfloor \sqrt{n} - 1 \rfloor$ ancillae, add it to a second register of $\lfloor \log_2 n \rfloor + 1$ qubits, and then uncompute the Hamming weight of the group. The same $\lfloor \sqrt{n} - 1 \rfloor$ ancillae are used to compute the Hamming weight of each group. After the Hamming weight of all $\lceil n/\lfloor \sqrt{n} \rfloor \rceil$ groups has been computed and added to the total, we apply phases to the total Hamming weight as before. Then, we recompute and subtract each of the group Hamming weights to uncompute the total Hamming weight register. This reduces the number of ancillae required from $n - 1$ as in the previous section, where the Hamming weight of all n initial rotation qubits was computed at once, to $\lfloor \sqrt{n} \rfloor + \lfloor \log_2 n \rfloor$.

The number of T gates required when we divide into groups of size $\lfloor \sqrt{n} \rfloor$, and add to a total Hamming weight register, is as follows. Computing the Hamming weight within each group can be done at a cost of at most $4 \lfloor \sqrt{n} - 1 \rfloor$ T gates. Adding that group Hamming weight to the total Hamming weight register costs at most $4 \lfloor \log_2 n \rfloor$ T gates. Uncomputing the group Hamming weight after adding it to the total register (to continue on to the next group) costs no T gates; however, to uncompute the total Hamming weight we must subtract the Hamming weight of each group. Because we uncompute the group Hamming weights to free up the $\lfloor \sqrt{n} \rfloor$ ancillae, this subtraction requires us to sequentially recompute and subtract each group Hamming weight. For the final group added to the total, we do not need to recompute the group Hamming weight because we can subtract immediately after phasing the total Hamming weight, without uncomputation.

Thus, we compute the Hamming weight of $\lceil n/\lfloor \sqrt{n} \rfloor \rceil - 1$ groups twice. The total cost of this is $8 \lceil n/\lfloor \sqrt{n} \rfloor \rceil - 1 \lfloor \sqrt{n} - 1 \rfloor$ T gates, and we compute the Hamming weight of the final group once at a cost of $4 \lfloor \sqrt{n} - 1 \rfloor$ T gates. We add and subtract the Hamming weight of each group to the total Hamming weight register once, at a total cost of $8 \lfloor \log_2 n \rfloor \lceil n/\lfloor \sqrt{n} \rfloor \rceil$ T gates. All of these values are upper bounds on the cost. The full T-cost of the $\lfloor \sqrt{n} \rfloor$ -grouping method is then

$$8 \left(\lceil n/\lfloor \sqrt{n} \rfloor \rceil - \frac{1}{2} \right) \lfloor \sqrt{n} - 1 \rfloor + 8 \lfloor \log_2 n \rfloor \lceil n/\lfloor \sqrt{n} \rfloor \rceil. \quad (16)$$

We can get a better understanding of this by simplifying using the inequality $\lceil n/\lfloor \sqrt{n} \rfloor \rceil < \sqrt{n} + 3/2$ for $n \in \mathbb{N}^{>0}$. This shows that the number of T gates required to reduce the original n rotations to $\lfloor \log_2 n \rfloor + 1$ rotations is (loosely) upper bounded by

$$8n + 8\sqrt{n} \log_2 n + 12 \log_2 n - 8. \quad (17)$$

The grouping method thus requires slightly more than double the $4n - 4$ T gates of the method described in the previous section. On the other hand, it only requires $\lfloor \sqrt{n} \rfloor + \lfloor \log_2 n \rfloor$ ancilla qubits, compared with the n required for that method.

We can use the different methods presented in this section within the same circuit, depending on how many rotations we wish to combine at a given point. Say we have at most $n = 500$ arbitrary repeated rotations to combine in a circuit. Then we assign $\lfloor \sqrt{n} \rfloor + \lfloor \log_2 n \rfloor = 30$ ancilla qubits for combining those 500 rotations via Hamming weight phasing. If at some other point in the circuit we have 50 rotations to combine, we can do so using the existing 30 ancillae and $4r \lceil n/(r+1) \rceil = 4 \times 24 \times 2 = 192$ T gates, reducing the number of arbitrary rotations to be synthesized from 50 to 10. Further improvements in the number of ancilla qubits required are possible using other grouping strategies, but are beyond the scope of this work. We refer the interested reader to work on pebble games (see, e.g., Chapter 10 of [86], or [87]).

A.3 Catalyzing two \sqrt{T} gates using 5 T gates

In this subsection, we present an application of Hamming weight phasing discussed in [81]. We mentioned in Appendix A.1 that when the rotation angles are $\theta = m\pi/2^k$ for integers k, m , it is most efficient to only combine until the rotations are T gates. These rotations appear in some of the circuits in the paper (specifically, in the FFFT for side lengths which are a power of two), and we present a circuit tailored to \sqrt{T} and $T\sqrt{T}$ gates in this section.

This circuit uses Hamming weight phasing with an ancilla qubit to catalyze \sqrt{T} gates on two other qubits using only one Toffoli and one T gate, or 5 T gates, given the seed state $\sqrt{T}|+\rangle$. The same circuit can catalyze two $\sqrt{T}^3 = T\sqrt{T}$ gates for the same fault-tolerant cost by changing the seed state and applying an S gate on the ancilla qubit in addition to the T gate. The circuit is constructed by moving the \sqrt{T} gate of the standard adder-based Hamming weight phasing circuit through to the beginning to create the seed state $\sqrt{T}|+\rangle$. The seed state must be synthesized at the beginning of the entire computation at full cost, but can be used anywhere after that point—it is not consumed as we generate \sqrt{T} gates. We show the catalysis circuit and its adder-based counterpart in Figure 4.

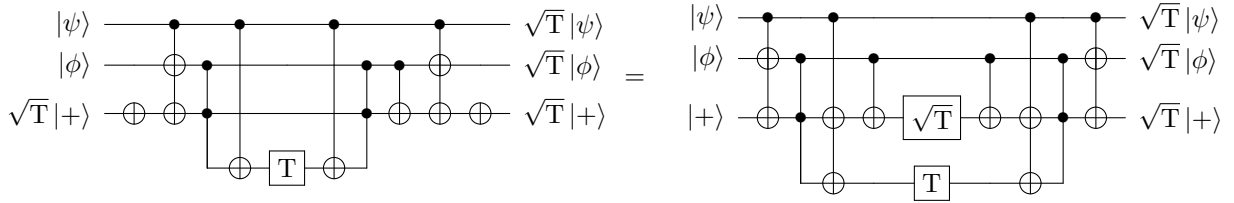


Figure 4: Circuit to catalyze two \sqrt{T} gates using 5 T gates (left), and equivalent adder-based circuit (right). The left circuit is constructed from the right by moving the \sqrt{T} gate through to the beginning using circuit identities. While the right circuit applies three \sqrt{T} gates using a \sqrt{T} and 5 T gates, the left circuit takes one of the \sqrt{T} gates as input in the form of a seed state, and applies the other two \sqrt{T} gates using 5 T gates. We use the “cornering” notation for temporary logical AND of [61] (see Figure 5). The same circuit catalyzes two $\sqrt{T}^3 = T\sqrt{T}$ gates if we replace the seed state with $\sqrt{T}^3|+\rangle$ and add an S gate on the ancilla qubit after the T gate.

We use the “cornering” notation of [61, 81] for computing and uncomputing the logical AND to indicate that the ancilla qubit is free for use elsewhere in the circuit outside that part of the computation. We review this notation in Figure 5 and present equivalent Clifford+T circuits.

B Trotter steps by fermionic swap network

We review in this appendix the fermionic swap network Trotter step developed in [29]. The idea for this Trotter step is that the Hamiltonian Eq. (1) can be simulated in linear depth in a network which changes the Jordan-Wigner ordering as it runs, such that the qubits corresponding to each spin-orbital are adjacent to the qubits corresponding to each other spin-orbital in the system at some point in the network.

The fermionic swap network does this by reversing the initial Jordan-Wigner ordering of the spin-orbitals with fermionic swaps using odd-even transposition sort (parallel bubble sort), simulating terms in the Hamiltonian that act on neighboring qubits as it goes. The sorting algorithm alternates between two different layers of swaps. In the first layer, each qubit in an odd-numbered position is fermionically swapped with the even-numbered qubit to its right. In the second layer, each qubit in an even-numbered position is fermionically swapped with the odd-numbered qubit to its right. This is shown for 5 qubits in the leftmost column of Table 3. After the final set of swaps, the original Jordan-Wigner ordering has been reversed: because this reversal was performed using only nearest-neighbor operations, for each term in the Hamiltonian there must have been a point in the network where the spin-orbitals the term acts on were mapped to adjacent qubits.

Now, suppose in a particular layer of the swap network that spin-orbital p (encoded by qubit i_p in the Jordan-Wigner ordering) undergoes a fermionic swap with spin-orbital q (encoded by qubit i_q). Since all fermionic swaps occur between adjacent qubits, either $i_q = i_p + 1$ or $i_p = i_q + 1$. Thus, the fermionic operator $V_{pq}n_p n_q$ can be simulated by the qubit operators $V_{pq}(\mathbb{1} - Z_{i_p} - Z_{i_p+1} + Z_{i_p} Z_{i_p+1})/4$, and the fermionic operators $T_{pq}(a_p^\dagger a_q + a_q^\dagger a_p)$ can be simulated by the qubit operators $T_{pq}(X_{i_p} X_{i_p+1} + Y_{i_p} Y_{i_p+1})/2$.

This shows how to simulate all the terms in the Hamiltonian involving spin-orbitals p and q . We integrate these evolutions with the fermionic swap, and refer to the combined unitary as the “fermionic simulation gate” following [29]. The fermionic simulation gate performs three operations: it simulates evolution for time t under $T_{pq}(X_{i_p} X_{i_p+1} + Y_{i_p} Y_{i_p+1})/2$, simulates evolution for time t under $V_{pq}(Z_{i_p} Z_{i_p+1} - Z_{i_p} - Z_{i_p+1} + \mathbb{1})/4$, and finally fermionically swaps the two spin-orbitals. In matrix form, this gate is

$$\mathcal{F}_t(i_p, i_q) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -i \sin(T_{pq}t) & \cos(T_{pq}t) & 0 \\ 0 & \cos(T_{pq}t) & -i \sin(T_{pq}t) & 0 \\ 0 & 0 & 0 & -e^{-iV_{pq}t} \end{pmatrix}. \quad (18)$$

We can see then that Table 3 actually depicts an entire first-order Trotter step if we interpret the boxes in the canonical ordering as the unitary $\mathcal{F}_t(i_p, i_q)$. If they are nonzero, the external potential terms $U_p n_p$ can be simulated by applying single-qubit rotations after any layer of fermionic simulation gates.

The second-order (symmetric) Trotter step can be constructed from the Trotter step described so far by running it once forward, and symmetrizing by running it again in reverse. The spin-orbitals are restored to their original ordering by the second-order Trotter step.

This Trotter step can simulate any Hamiltonian of the form Eq. (1), including the Fermi-Hubbard model, the uniform electron gas, and the electronic structure problem. For the Fermi-Hubbard model on an $n \times n$ lattice without periodic boundary conditions, we only need to run through the first $\mathcal{O}(\sqrt{n})$ layers of fermionic simulation gates [29]. With periodic boundary conditions, the full network must be

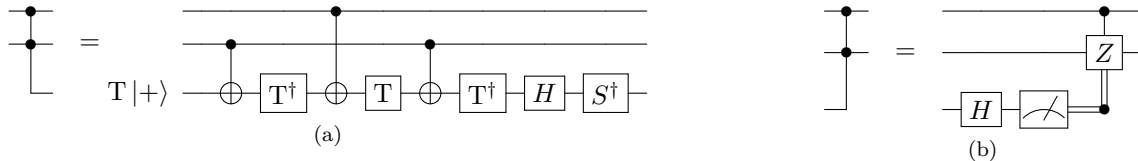


Figure 5: Circuits with cornering notation for computing and uncomputing logical AND operations, together with the equivalent circuits in the Clifford+T basis [61]. The cornering notation indicates that the qubit is only needed during that section of the circuit. In Figure 4, the ancilla qubit is available for use before computing and after uncomputing the logical AND. (a) Computing the logical AND requires one Toffoli or 4 T gates; (b) uncomputing it requires none.

Canonical Ordering	$\frac{1}{2} (X_1 X_2 + Y_1 Y_2)$	$\frac{1}{2} (X_2 X_3 + Y_2 Y_3)$	$\frac{1}{2} (X_3 X_4 + Y_3 Y_4)$	$\frac{1}{2} (X_4 X_5 + Y_4 Y_5)$
$\boxed{\varphi_1 \varphi_2} \boxed{\varphi_3 \varphi_4} \varphi_5$	$a_1^\dagger a_2 + a_2^\dagger a_1$		$a_3^\dagger a_4 + a_4^\dagger a_3$	
$\varphi_2 \boxed{\varphi_1 \varphi_4} \boxed{\varphi_3 \varphi_5}$		$a_1^\dagger a_4 + a_4^\dagger a_1$		$a_3^\dagger a_5 + a_5^\dagger a_3$
$\boxed{\varphi_2 \varphi_4} \boxed{\varphi_1 \varphi_5} \varphi_3$	$a_2^\dagger a_4 + a_4^\dagger a_2$		$a_1^\dagger a_5 + a_5^\dagger a_1$	
$\varphi_4 \boxed{\varphi_2 \varphi_5} \boxed{\varphi_1 \varphi_3}$		$a_2^\dagger a_5 + a_5^\dagger a_2$		$a_1^\dagger a_3 + a_3^\dagger a_1$
$\boxed{\varphi_4 \varphi_5} \boxed{\varphi_2 \varphi_3} \varphi_1$	$a_4^\dagger a_5 + a_5^\dagger a_4$		$a_2^\dagger a_3 + a_3^\dagger a_2$	
$\varphi_5 \varphi_4 \varphi_3 \varphi_2 \varphi_1$				

Table 3: The five layers of swaps for the five spin-orbital fermionic swap network. The boxes in the leftmost column of each row indicate the pairs of qubits involved in a fermionic simulation gate in a given layer. The right columns show which kinetic energy fermion operators the different nearest-neighbor qubit operators correspond to at that layer. The final ordering is shown in the last row of the table for completeness. All $\binom{5}{2} = 10$ possible kinetic energy interactions appear at some point in the table, and are thus applied during the steps shown. The fermionic simulation gate also applies the potential energy interactions $V_{pq} n_p n_q$ at that time. Thus, these five layers implement an entire Trotter step of the Hamiltonian. A slash through a cell indicates that no interaction is performed between that pair of qubits in the layer.

run to ensure that interactions between the top and bottom rows of the lattice are simulated. For the electronic structure problem, the external potential terms $U_p n_p$ arise from the charges of the nuclei ζ_j . These nuclear charges are all that distinguish various molecules and materials from the uniform electron gas. However, because they contribute only to the terms $U_p n_p$, they enter the simulation only as a layer of single-qubit rotations appended to the Trotter step circuit.

C Trotter steps of the split-operator algorithm

In this appendix, we review how to simulate Trotter steps of the Hamiltonian Eq. (1) using a variation of the split-operator algorithm introduced in [28]. The general principle of a split-operator algorithm is to divide (“split”) the Hamiltonian into parts which can be more easily simulated separately. In this case, the algorithm separately simulates the kinetic energy operator in the momentum basis, and the potential energy operator in the position basis. We use an additional circuit to change back and forth between the position and momentum bases. In the momentum basis, the kinetic energy operator is diagonal, and moreover can be simulated using only single-qubit rotations. In the position basis, the potential energy operator is diagonal, containing only terms of the form $Z_p Z_q$. While this can be simulated using a swap network as in [28] or [29], we will instead develop a specialized network tailored to take full advantage of Hamming weight phasing.

We begin this appendix by describing how to construct the split-operator Trotter step. We then introduce our specialized network for taking full advantage of Hamming weight phasing when simulating the potential energy terms. Following that, we present the circuits used in the split-operator algorithm to change between the position and momentum bases. When the system side length is a power of 2, we use the fermionic fast Fourier transform (FFFT) as in [28, 42]. We compute the T-cost of the FFFT using rotation catalysis as discussed in Appendix A.3. For other system side lengths, we use an alternative approach to diagonalizing the kinetic energy operator using Givens rotations [29, 43]. While this second approach has fewer restrictions on the kinetic energy operators it can diagonalize than the FFFT, in the fault-tolerant setting, it is more costly because it uses more arbitrary rotations.

C.1 Constructing the split-operator Trotter step

We discuss here how to construct the split-operator Trotter step. Specifically, we will describe the circuit for the second-order (first-order symmetric) split-operator Trotter step.

Recall from Eq. (1) that we wish to simulate evolution under a fermionic Hamiltonian of the form

$$H = \sum_{pq} T_{pq} a_p^\dagger a_q + \sum_p U_p n_p + \sum_{p \neq q} V_{pq} n_p n_q. \quad (19)$$

In the context of the electronic structure problem, T_{pq} , U_p , and V_{pq} are real numbers defined by integrals over molecular orbitals, arising from the kinetic energy, external potential, and electron-electron Coulomb repulsion, respectively. The algorithms we discuss can simulate arbitrary Hamiltonians of this form, not necessarily arising from any physical system. The creation and annihilation operators index the spin-orbitals through vectors p and q . This second-quantized Hamiltonian can be mapped by the Jordan-Wigner transform to a qubit Hamiltonian of the form

$$H = \sum_{pq} \tilde{T}_{pq} (X_p Z_{r_i} Z_{r_{i+1}} \cdots Z_{r_{i+k}} X_q + Y_p Z_{r_i} Z_{r_{i+1}} \cdots Z_{r_{i+k-1}} Z_{r_{i+k}} Y_q) + \sum_p \tilde{U}_p Z_p + \sum_{pq} \tilde{V}_{pq} Z_p Z_q, \quad (20)$$

where we have omitted the identity term because it is trivial to simulate. Here, X_p , Y_p , and Z_p are the Pauli X , Y , and Z gates, respectively, acting on the qubit indexed by p in the Jordan-Wigner ordering. $Z_{r_i} \cdots Z_{r_{i+k}}$ is the string of Pauli Z gates between qubit p and qubit q in the Jordan-Wigner order, needed to preserve the appropriate commutation relations for the transformed fermionic creation and annihilation operators. The coefficients \tilde{T}_{pq} , \tilde{U}_p , and \tilde{V}_{pq} in the qubit Hamiltonian can be directly computed from the coefficients without tildes in the second-quantized Hamiltonian.

The split-operator Trotter step splits evolution e^{-iHt} under the Hamiltonian H into separate evolutions under the kinetic and potential energy parts of the Hamiltonian. The reason this is useful is that the terms in the kinetic energy part of the Hamiltonian (those with coefficients \tilde{T}_{pq}) can be diagonalized by an efficient circuit transformation C , leading to a sum of single-qubit terms, that is,

$$\sum_{pq} \tilde{T}_{pq} (X_p Z_{r_i} Z_{r_{i+1}} \cdots Z_{r_{i+k}} X_q + Y_p Z_{r_i} Z_{r_{i+1}} \cdots Z_{r_{i+k-1}} Z_{r_{i+k}} Y_q) = C \left(\sum_p T_p Z_p \right) C^\dagger. \quad (21)$$

The diagonalizing circuits C and C^\dagger change between the position and momentum bases using either the FFFT or Givens rotations. The idea of the split-operator method is that the Trotter-Suzuki approximation can be applied to divide evolution under the total Hamiltonian H into two evolutions: one under the kinetic energy terms

$$T = C \left(\sum_p T_p Z_p \right) C^\dagger, \quad (22)$$

that is e^{-iTt} , and one under the potential energy terms

$$U + V = \sum_p \tilde{U}_p Z_p + \sum_{pq} \tilde{V}_{pq} Z_p Z_q, \quad (23)$$

that is $e^{-i(U+V)t}$. Symmetrizing these evolutions to give the second-order Trotter step, we approximate time evolution under the full Hamiltonian by blocks of evolutions separated by the circuit C and its inverse C^\dagger ,

$$e^{-iHt} \approx C \left[\prod_p e^{-iT_p Z_p t/2} \right] C^\dagger \left[\prod_p e^{-i\tilde{U}_p Z_p t} \prod_{pq} e^{-i\tilde{V}_{pq} Z_p Z_q t} \right] C \left[\prod_p e^{-iT_p Z_p t/2} \right] C^\dagger. \quad (24)$$

The equation above describes a single second-order Trotter step: the approximation can be refined by dividing into r such steps, each of time t/r .

In fact, we have two choices for how to symmetrize the split-operator Trotter step: as opposed to Eq. (24), in which we approximate evolution for r Trotter steps as $e^{-iHt} \approx (e^{-iTt/2r} e^{-i(U+V)t/r} e^{-iTt/2r})^r$, we can instead approximate it as $e^{-iHt} \approx (e^{-i(U+V)t/2r} e^{-iTt/r} e^{-i(U+V)t/2r})^r$. Including the circuits for diagonalizing the kinetic energy operator, a single Trotter step of this form is

$$e^{-iHt} \approx \left[\prod_p e^{-i\tilde{U}_p Z_p t/2} \prod_{pq} e^{-i\tilde{V}_{pq} Z_p Z_q t} \right] C \left[\prod_p e^{-iT_p Z_p t} \right] C^\dagger \left[\prod_p e^{-i\tilde{U}_p Z_p t} \prod_{pq} e^{-i\tilde{V}_{pq} Z_p Z_q t/2} \right]. \quad (25)$$

At first glance, the number of gates required for the two Trotter steps might appear to be different. The number of terms in T is linear in the number of spin-orbitals, whereas the number of terms in $U + V$ is quadratic in the number of spin-orbitals, which would lead us to expect that Eq. (25)—in which $e^{-i(U+V)t}$ appears twice—might require many more gates than Eq. (24), in which $e^{-i(U+V)t}$ only appears once. On the other hand, C or C^\dagger appear four times in Eq. (24) but only twice in Eq. (25), which might lead us to believe that Eq. (24) could use more gates.

However, this intuition is deceptive: so long as the number of Trotter steps r is large, since the beginning and end of both split-operator steps are the same and hence can be merged together, the difference in the costs is negligible. The result of this merging is that the cost of a simulation with r Trotter steps is nearly identical for the two competing orderings of terms. Only in the final Trotter step do the two orderings differ in cost: Eq. (24) applies C one more time and $e^{-i(U+V)t}$ one fewer time than Eq. (25). Both variants apply C , C^\dagger , e^{-iTt} , and $e^{-i(U+V)t}$ once in each of the remaining $r - 1$ Trotter steps, far outweighing this final difference.

Rather, it is differences in the number of Trotter steps r required for accurate simulation—dictated by the Trotter error of the two steps—that dominate the difference in cost. For all systems considered, we construct the Trotter step depending on which split-operator step yields the smaller Trotter error. For the system parameters chosen, simulating V followed by T yields a smaller Trotter error norm for the uniform electron gas with Wigner-Seitz radius $r_s \gtrsim 1$, whereas T followed by V yields a smaller Trotter error for the Hubbard model and the uniform electron gas with Wigner-Seitz radius $r_s \lesssim 1$. There is little significance to this beyond the relative sizes of the norms $\|T\|$ and $\|V\|$ (see Eq. (60), Eq. (61), and the following discussion).

Implementing the different circuits in the split-operator steps is straightforward. Each term $e^{-iT_p Z_p t}$, $e^{-i\tilde{U}_p Z_p t}$, or $e^{-i\tilde{V}_{pq} Z_p Z_q t}$ requires one arbitrary rotation, whether we implement time evolution, or the evolution directionally controlled on an ancilla $|c\rangle |\psi\rangle \mapsto |c\rangle e^{-iHt(-1)^c} |\psi\rangle$ that we shall use for phase estimation (Appendix F). Thus, with N spin-orbitals, at most $N(N + 3)/2$ arbitrary rotations are needed to simulate all the terms in the Hamiltonian, and the Trotter step circuit can be completed with one application each of the basis change circuit C and its inverse C^\dagger . We address two remaining questions in this appendix: first, how to order the simulation of the terms in $\prod_{pq} e^{-i\tilde{V}_{pq} Z_p Z_q t}$ so as to take full advantage of Hamming weight phasing, thereby reducing the number of arbitrary rotations required; and second, how to construct the circuit C using either the FFFT or Givens rotations.

C.2 Reducing the cost of simulating the potential energy operator using Hamming weight phasing

In this subsection we present a technique for reducing the number of arbitrary rotations needed to implement evolution under the interaction part of the potential energy operator $\prod_{pq} e^{-i\tilde{V}_{pq} Z_p Z_q t}$ in the split-operator Trotter step. Evolution under the terms $\prod_p e^{-i\tilde{U}_p Z_p t}$ can be done using only a single arbitrary rotation on each qubit. By comparison, $\prod_{pq} e^{-i\tilde{V}_{pq} Z_p Z_q t}$ in general appears to require $\mathcal{O}(N^2)$ arbitrary rotations for N qubits. We will show here how to use Hamming weight phasing to reduce this cost for the electronic structure problem to $\mathcal{O}(N \log N)$ arbitrary rotations, at a cost of $(N - 1)(N - 2)$

additional Toffoli or $4(N-1)(N-2)$ additional T gates. Because the number of T gates needed to synthesize each arbitrary rotation is much larger than 4, this greatly reduces the fault-tolerant cost of simulating these terms.

The idea behind our technique is that in the basis of [28], the coefficients $\tilde{V}_{p,q}$ are translation-invariant for the electronic structure Hamiltonian, that is, for any index vector s ,

$$\tilde{V}_{p,q} = \tilde{V}_{p+s,q+s}. \quad (26)$$

We will harness this property so that we can apply Hamming weight phasing to groups of angles that are as large as possible. Hamming weight phasing allows M parallel rotations by the same angle to be reduced to $\lceil \log_2 M + 1 \rceil$ rotations by different angles, and $4M - 4$ additional T gates (alternatively, $M - 1$ additional Toffoli gates), using $M - 1$ ancilla qubits. We can fully make use of Hamming weight phasing if we can order the terms $e^{-i\tilde{V}_{pq}Z_pZ_qt}$ in the simulation such that equiangular rotations are grouped together as much as possible. Note that we can compute and uncompute the parity of qubits p and q at no additional cost using a CNOT gate.

We do this by performing the evolution $\prod_{pq} e^{-i\tilde{V}_{pq}Z_pZ_qt}$ in layers such that the value of every \tilde{V}_{pq} simulated within a given layer is the same. In this case, the arbitrary rotations used in each layer are all by the same angle, allowing them to be fully combined using Hamming weight phasing. We accomplish this as follows. We choose the interactions simulated in each layer according to the set of pairs of qubits $(p, p+s)$ with fixed s , with the index vector p running over all qubits. The translation-invariance of $\tilde{V}_{p,q}$ guarantees, for all qubit indices p , that

$$\tilde{V}_{0,s} = \tilde{V}_{p,p+s}. \quad (27)$$

We first consider the case where the number of qubits N is even; when N is odd the layers differ only slightly from the even N case. For N qubits (with N even), there are N such pairs of qubits $(p, p+s)$, given by the $N-1$ different vectors s on the grid. However, because each of these interactions is between two qubits, at most $N/2$ of them can be performed simultaneously in a layer. We show a diagram of the two layers used for $s = (0, 1)$ in Figure 6 for the case of 16 qubits on a 4×4 grid. Note that the vector indices on the qubits are evaluated modulo 4 for the 4×4 grid. In each layer, we are able to simulate $N/2$ interactions of equal strength, yielding $N/2$ equiangular rotations combinable using Hamming weight phasing.

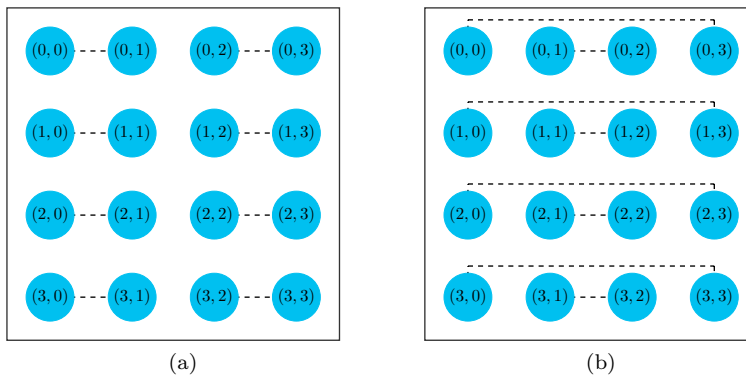


Figure 6: Two-qubit evolution terms (dotted lines) simulated in the two layers of the swap network with the index vector $s = (0, 1)$. In both layers, qubit p interacts with qubit $p + s$, where the elements of the index vectors p and $p + s$ are evaluated modulo the grid length. (a) Qubit p interacts with qubit $p + s$, starting with $p = (0, y)$ and $p = (2, y)$. (b) Qubit p interacts with qubit $p + s$, starting with $p = (1, y)$ and $p = (3, y)$.

In general, there are $N - 1$ different index vectors s that must be iterated over. (The final vector, $s = 0$, does not appear because it corresponds to a qubit interacting with itself.) Each index vector is

applied in two layers, leading to $2N - 2$ layers each simulating $N/2$ interactions. In each of the $2N - 2$ layers, Hamming weight phasing can be applied over every simulated term $e^{-i\tilde{V}_{p,p+s}Z_{p+s}Z_s t}$, because each coefficient $\tilde{V}_{p,p+s}$ is the same, and hence each of the $N/2$ rotation angles is the same.

This allows Hamming weight phasing to be applied over every rotation in each layer, reducing the number of arbitrary rotations required per layer from $N/2$ to $\lceil \log_2(N/2) + 1 \rceil$, at a cost of an additional $2N - 4$ T gates or $N/2 - 1$ Toffoli gates. Over all the $2N - 2$ layers, the T-cost to simulate all the two-qubit terms in a step of Trotter evolution is

$$4(N - 1)(N - 2) + T_{\text{synth}}(2N - 2)\lceil \log_2 N \rceil, \quad (28)$$

where $T_{\text{synth}} = \mathcal{O}(\log(1/\epsilon))$ is the number of T gates to synthesize each arbitrary rotation, with ϵ the required synthesis precision.

When the number of qubits N is odd, we instead must simulate each of the $N - 1$ index vectors using three rather than two layers. Because the number of qubits is odd, one qubit is left out in each layer as compared to the even case—its partner is already matched with another qubit. As a result there are $2N - 2$ layers of $(N - 1)/2$ equiangular rotations and $N - 1$ layers with a lone rotation in the simulation. We use Hamming weight phasing to reduce the $(N - 1)/2$ equiangular rotations per layer to $\lceil \log_2(N - 1)/2 + 1 \rceil$ at a cost of an additional $(N - 1)/2$ Toffoli or $2N - 2$ T gates. Over all the layers, the T-cost to simulate all the two-qubit terms in a step of Trotter evolution is

$$4(N - 1)^2 + T_{\text{synth}}(N - 1)(2\lceil \log_2(N - 1) \rceil + 1), \quad (29)$$

for the case of odd N .

For both even and odd N , we reduce the cost of the arbitrary rotations to $\mathcal{O}(N \log N \log(1/\epsilon))$ using at most $4(N - 1)^2$ additional T gates or $(N - 1)^2$ Toffoli gates. This is to be compared with the $\mathcal{O}(N^2 \log(1/\epsilon))$ cost were we to not apply Hamming weight phasing, or were the interaction terms not translation-invariant. This improvement applies to any quantum simulation of electronic structure using the basis of [28], or more generally to any simulation of a two-qubit interaction with the translation-invariance property Eq. (26).

C.3 The fast fermionic Fourier transform

The fermionic fast Fourier transform (FFFT) is one way of implementing the circuit C which diagonalizes the kinetic energy operator. In this appendix, we define the FFFT, present circuits for it, and compute the T-cost of using it to change from the position to momentum bases using rotation catalysis.

The effect of the fermionic Fourier transform is a single-particle rotation

$$c_\nu^\dagger = \frac{1}{\sqrt{N}} \sum_p a_p^\dagger e^{-ik_\nu \cdot r_p} \quad c_\nu = \frac{1}{\sqrt{N}} \sum_p a_p e^{ik_\nu \cdot r_p} \quad (30)$$

with $r_p = p(2\Omega/N)^{1/d}$ and $k_\nu = 2\pi\nu/\Omega^{1/d}$ the centroid of the orbital given by the index vector p , and k_ν a momentum mode as in Section 2.

The circuit for the FFFT is composed of two types of local gates. The first is the fermionic swap gate, which has the effect of exchanging two orbitals, while maintaining proper anti-symmetrization. Under the Jordan-Wigner transform, the fermionic swap gate f_{swap} has the matrix representation

$$f_{\text{swap}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad (31)$$

i.e. it swaps two qubits and applies a sign -1 to the $|11\rangle$ state. The second gate used for the FFFT has the matrix representation

$$F_{k,n} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{\exp(2\pi ik/n)}{\sqrt{2}} & 0 \\ 0 & \frac{1}{\sqrt{2}} & -\frac{\exp(2\pi ik/n)}{\sqrt{2}} & 0 \\ 0 & 0 & 0 & -\exp(2\pi ik/n) \end{pmatrix} \quad (32)$$

where n is the number of qubits the FFFT acts on, and k is an integer. This has the effect of a rotation on the first qubit followed by a Hadamard transform in the single-particle subspace. Our notation for this gate is slightly different from other works discussing the FFFT [28, 42] which denote the same gate by F_k (with the number of qubits n left implicit). In this sense our presentation is closer to that in [72].

The FFFT on n qubits can be implemented recursively in depth $\mathcal{O}(n \log n)$ using only the fermionic swap gate f_{swap} and the gates $F_{k,n}$ for integer k [28]. This recursive structure constrains the FFFT to operate on systems whose side lengths are powers of 2. The strategy for this recursive circuit parallels the recursive structure of the classical fast Fourier transform, and can be broken into four stages.

1. First, the $n/2$ -qubit FFFT is applied separately to the first and second halves of the qubit array. The base case is the two-qubit FFFT, which is the gate $F_{0,2}$.
2. Next, an interleaving step of fermionic swaps puts the qubits in bit-reversed order. (The qubits are re-ordered based on their index, viewed as binary digits, in reverse. For example $[0, 1, 2, 3]$ written in binary digits is $[00, 01, 10, 11]_b$ and hence would be re-ordered as $[00, 10, 01, 11]_b = [0, 2, 1, 3]$.) The purpose of this reordering step is to separate the groups of qubits by parity of the index—they are now separated by their final digit. The recursive substructures of the FFFT operate differently depending on whether the final digit is 0 or 1.
3. The reordered qubits are paired, and $F_{i,n}$ applied to phase the i th pair for $i \in [0, n/2 - 1]$.
4. Finally, a de-interleaving step restores the qubits to their initial order.

We present the circuit for the 8-qubit FFFT on a nearest neighbor architecture in Figure 7, emphasizing the different stages in the circuit. Stage 3 is only a single layer of gates and we merge it into stage 2.

So far, we have discussed the structure of the FFFT circuit. Next we will compute the number of T gates required to implement it to determine the cost of the FFFT on a fault-tolerant quantum computer. We do this by presenting explicit circuits for the gate $F_{k,n}$, and counting how many of them appear in the circuit.

The recursive structure of the n -qubit FFFT (where n is necessarily a power of two) means that it uses $n/2$ $F_{k,n}$ gates in $\log_2 n$ recursive layers, yielding $(n \log_2 n)/2$ $F_{k,n}$ gates in total. The second gate used in the circuit, the fermionic swap gate, appears $n(n-2)/4$ times directly in the interleave and de-interleave steps of the n -qubit FFFT, as well as in the interleave/de-interleave steps of the two recursive $n/2$ -qubit FFFTs. In total, this is

$$\sum_{i=1}^{\log_2 n - 1} n(n - 2^{i-1})/2^i = n(n - \log_2(8n)/2) \quad (33)$$

f_{swap} gates. In fact, f_{swap} gates are free in this model, requiring only Clifford gates, and we will not discuss them further here. By comparison, any fault-tolerant circuit construction for $F_{k,n}$ will require T gates. We present one such construction in Figure 8—note that $T^\dagger = ZST$. $F_{k,n}$ requires two T gates when $4k/n$ is an integer, three T gates if instead $8k/n$ is an integer, and two T gates and an arbitrary rotation otherwise. In the particular case that $16k/n$ is an integer, the arbitrary rotation is a

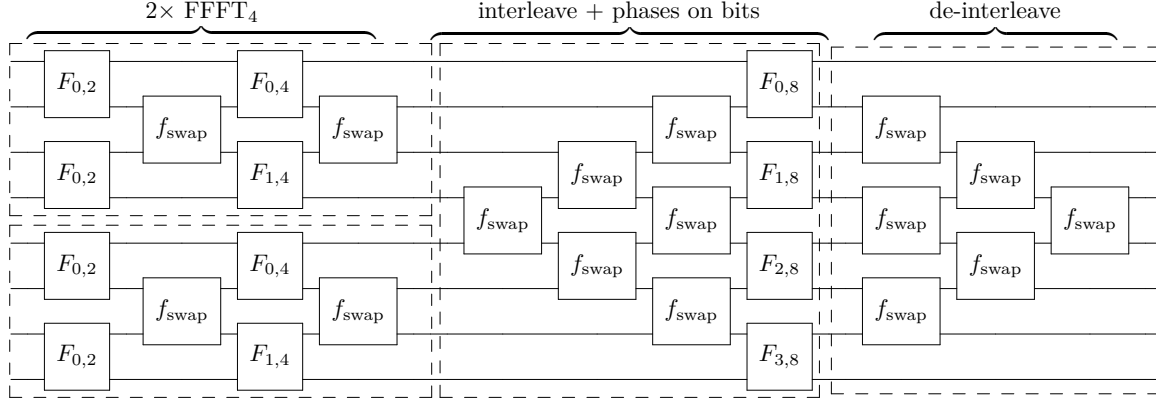


Figure 7: Circuit for the 8-qubit fermionic fast Fourier transform (FFFT₈) using only two-qubit nearest-neighbor gates. The dotted boxes denote the different stages of the decomposition of the FFFT circuit. Furthest left are the two 4-qubit FFFTs (FFFT₄) that appear in the recursive structure of the 8-qubit FFFT. Within the 4-qubit FFFT, the two-qubit gate $F_{0,2}$ is the recursive base case: the 2-qubit FFFT is exactly the gate $F_{0,2}$. The triangular structures of fermionic swap gates interleave the qubits (putting them in bit-reversed order, as described in the text) before they are phased by the gates $F_{i,8}$. Finally, the de-interleaving step returns the qubits to their original ordering. Note that several of the phasing gates in the circuit are the same: $F_{k,n} = F_{j,m}$ if $k/n = j/m$. For example, $F_{1,4}$ in the 4-qubit gate is equivalent to $F_{2,8}$ in the 8-qubit FFFT.

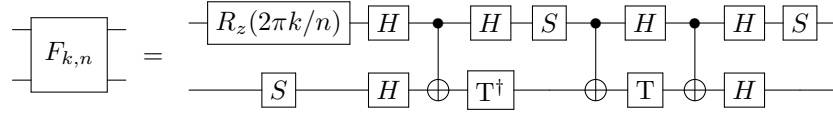


Figure 8: Fault-tolerant circuit for the two-qubit fermionic Fourier transform gate $F_{k,n}$ in the n -mode FFFT, up to global phase. The T and S gates are $T = R_z(\pi/4)$ and $S = T^2$. If $4k/n$ is an integer, the rotation at the beginning is at most an S gate; if instead $8k/n$ is an integer, it is an additional (third) T gate. The largest case we consider is applying the FFFT to a system with side length $n = 16$: in that case, \sqrt{T} gates can appear in the rotation. These can be synthesized in pairs using only 5 T gates per pair, rather than the cost of a full arbitrary rotation.

\sqrt{T} gate. Recall from Appendix A.3 that provided that \sqrt{T} gates appear two at a time in the circuit, we can implement them using only 2.5 T gates each.

This means that the T-cost of the FFFT circuit is 8 for the 4-qubit case, 26 for the 8-qubit case, and 71 T gates + 4 \sqrt{T} gates = 81 T gates for the 16-qubit case, not accounting for the cost of preparing the seed $\sqrt{T}|+\rangle$ state. Next, we will shift our attention to the FFFT in multiple dimensions. Following that, we will show how to minimize the T-cost of the 2D FFFT using rotation catalysis, when the system being Fourier transformed is large enough that the FFFT has non-trivial rotations (rotations which must be synthesized).

C.3.1 Changing basis in multiple dimensions

We use the fermionic Fourier transform or Givens rotations to implement the circuit C to diagonalize the kinetic energy terms in the Hamiltonian in the split-operator algorithm. Earlier in this section, we discussed the number of arbitrary rotations required for the fast fermionic Fourier transform on n qubits in one dimension. For a d -dimensional system, however, the 1D FFFT circuit must be applied along each of the system dimensions. Though we have not yet discussed it, this also holds if we implement C using Givens rotations, in the sense that whether we use the FFFT or Givens rotations,

the d -dimensional basis change is composed of the same pattern of 1D basis changes. The discussion in this section applies identically to both approaches.

In each dimension, the 1D basis change must be applied once for each value of the other coordinates. For example, for a hypercube with side length n in d dimensions, transforming along one dimension requires n^{d-1} applications of the 1D basis change. Repeating this once for each dimension, the 1D basis changing circuit must be applied dn^{d-1} times, in d layers of n^{d-1} basis changing circuits.

More concretely, for a $4 \times 4 \times 4$ grid, we apply the 4-qubit FFFT 16 times to Fourier transform the x dimension. The 1D 4-qubit FFFT is applied to the 4 qubits with fixed (y, z) and varying x coordinate. There are 16 pairs (y, z) so the 4-qubit FFFT is applied 16 times. This applies the fermionic Fourier transform in the x dimension: the procedure is repeated for the y dimension (to each of the 16 groups of 4 qubits with the same x and z coordinates, with the y coordinate varying), and again for the z dimension (to groups with the x and y coordinates now fixed, and the z coordinate varying). In total, the 4-qubit FFFT is applied $dn^{d-1} = 48$ times for this system.

For a system with spin, there are still d parallel stages in which the 1D basis changing circuit is applied multiple times. However, in each stage, the 1D circuit is applied $2n^{d-1}$ times, twice as many as in the corresponding spinless system. This doubling occurs because of the two spin sectors.

In general, this means that the 1D basis transformation circuits are applied many times in parallel, whether we use the FFFT or Givens rotations. Because the Givens rotation circuits contain arbitrary rotations, this allows us to apply Hamming weight phasing (Appendix A) to reduce the number of these arbitrary rotations which must be synthesized. For the FFFT, because the rotations are always of the form $R_z(2\pi k/n)$ with n a power of two, Hamming weight phasing is less efficient. Instead, as we discuss in the next section, we can use rotation catalysis as in Appendix A.3 to apply e.g. \sqrt{T} and $T\sqrt{T}$ gates using only T gates, without the usual costs of rotation synthesis. We discuss this in the next subsection for the case of the 16×16 grid, both with and without spin.

C.3.2 Rotation catalysis and the 2D FFFT

Of the systems we study, there is only one grid size for which rotation catalysis (see Appendix A.3) is beneficial for the FFFT. This is because the FFFT only has a non-trivial rotation—a rotation which is not entirely composed of T, S , and Z gates—for systems with side lengths greater than 8. Given that the FFFT can only be applied when the side length is a power of 2, the first such side length is 16. At this size, the FFFT circuit has several \sqrt{T} and $\sqrt{T}^3 = T\sqrt{T}$ gates. The 16×16 grid appears for the largest spinless uniform electron gas system we consider, and with spin for the Fermi-Hubbard model.

Fourier transforming with side length $n = 16$ in $d = 2$ dimensions applies the 16-qubit (1D) FFFT 32 times without spin, and 64 times with spin. This represents two layers each of 16 (32) parallel applications of the 1D circuit without (with) spin. For the 16-qubit FFFT, the rotation $R_z(2\pi k/n)$ in the gate $F_{k,16}$ is a $\{\sqrt{T}, T\sqrt{T}, S\sqrt{T}, ST\sqrt{T}\}$ gate for $k \in \{1, 3, 5, 7\}$, respectively. Ignoring the Clifford gates, the 16-qubit FFFT uses two \sqrt{T} and two $T\sqrt{T}$ gates. Recall from Figure 8 that every $F_{k,n}$ gate costs two T gates in addition to these, and that there are $(n \log_2 n)/2 = 32$ $F_{k,n}$ gates in the 16-qubit FFFT. Further, for $k = 2$ and $k = 6$, the additional rotation in $F_{k,16}$ contains a third T gate, and $F_{1,8}$ and $F_{3,8}$ contribute another 4 T gates (these gates appear twice, since the 8-qubit FFFT appears twice in the recursive circuit for the 16-qubit FFFT). There are thus 70 T gates in the 16-qubit FFFT.

Because the 16-qubit FFFT is applied 32 (64) times without (with) spin, changing between the position and momentum basis using it would appear to require us to distill 2240 (4480) T gates and additionally synthesize 128 (256) \sqrt{T} gates. Instead, by using rotation catalysis, we can reduce this to two arbitrary rotations to create the seed states at the beginning of the computation, regardless of how many split-operator Trotter steps we apply. We then use these seed states to catalyze all the \sqrt{T} and $\sqrt{T}^3 = T\sqrt{T}$ gates at a cost of half a Toffoli and half a T gate each, rather than the ~ 40 –50 T gates per rotation required for synthesis. The only downside is the need for three ancilla qubits:

two to store the seed states $\sqrt{T}|+\rangle$ and $T\sqrt{T}|+\rangle$, and one to temporarily compute logical ANDs as in Figure 4.

In total, converting between the position and momentum bases using the 2D fermionic Fourier transform on a 16×16 grid requires 2304 T and 64 Toffoli gates without spin and 4608 T and 128 Toffoli gates with spin. Because the other cases for which we use the FFFT (systems with side length 4 and 8) do not require arbitrary rotations in the FFFT, there is no advantage to these catalysis techniques. Finally, for side lengths that are not a power of two, we cannot construct the diagonalizing circuit C using the above FFFT construction. For these cases, we instead diagonalize using Givens rotations. We elaborate upon this diagonalization process in the next subsection.

C.4 Diagonalizing the kinetic energy operator using Givens rotations

Whereas the FFFT allows transformations of the fermionic mode operators based on the Fourier transform, Givens rotations allow arbitrary unitary transformations to be performed on the mode operators [29]

$$\tilde{a}_i^\dagger = \sum_j U_{ij} a_j^\dagger, \quad \tilde{a}_i = \sum_j U_{ij}^* a_j \quad (34)$$

for unitary U , with a_i^\dagger and a_i the initial creation and annihilation operators, respectively. This means we can construct a basis change circuit C capable of diagonalizing any number-conserving quadratic Hamiltonian $\sum_{pq} T_{pq} a_p^\dagger a_q$. By comparison, implementations of the diagonalizing circuit C using the FFFT are limited to the case where \tilde{a}_i^\dagger is the Fourier transform of a_i^\dagger ; the recursive structure (radix-2 decimation) further restricts the FFFT to act on a number of qubits n which is a power of 2. Other radices in the fermionic Fourier transform are possible in principle (and in fact their recursive base cases could be implemented using Givens rotations), but the corresponding circuits require more rotation gates than the radix-2 FFFT. Whereas the most challenging gates in the FFFT are rotation gates $R_z(2\pi k/n)$ with n a power of 2, as we shall see the majority of the rotation gates used when we construct C with Givens rotations are by arbitrary angles.

The reasoning behind the diagonalization of the kinetic energy operator using Givens rotations is as follows. By the Thouless theorem, applying the single-particle ($n \times n$) unitary transformation U is equivalent to applying the $2^n \times 2^n$ unitary

$$\exp\left(\sum_{ij} [\log U]_{ij} (a_i^\dagger a_j - a_j^\dagger a_i)\right). \quad (35)$$

This unitary can be implemented using a sequence of rotations of the form [29]

$$R_{ij}(\theta) = \exp[\theta_{ij}(a_i^\dagger a_j - a_j^\dagger a_i)]. \quad (36)$$

The action of $R_{ij}(\theta)$ is a Givens rotation by θ in the single-particle subspace. Appropriate choices of the angles θ allows us to diagonalize the matrix U , leaving n phases along the diagonal. The unitary U can be applied by applying these n phases to the qubit corresponding to that spin-orbital, followed by the generated sequence of Givens rotations in reverse. For n orbitals, $\binom{n}{2}$ Givens rotations are sufficient for each spin sector. The spinless case thus requires exactly $n(n-1)/2$ Givens rotations, whereas for the case with spin $n(n-1)$ Givens rotations are required [29].

We show a fault-tolerant circuit implementing $R_{ij}(\theta)$ between neighboring qubits in Figure 9. Each Givens rotation $R_{ij}(\theta)$ requires 2 arbitrary rotations, so the total number of arbitrary rotations required per spin sector is $n(n-1)$.

In total, without spin, the Givens rotations require $n(n-1)$ arbitrary rotations for any single-particle basis change using this procedure, while with spin $2n(n-1)$ are needed. The diagonal phases require n arbitrary rotations without spin and $2n$ arbitrary rotations with spin.

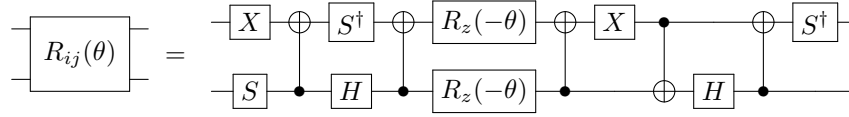


Figure 9: Fault-tolerant circuit for $R_{ij}(\theta) = \exp[\theta_{ij}(a_i^\dagger a_j - a_j^\dagger a_i)]$ on neighboring qubits, corresponding to a Givens rotation in the single-particle subspace. Two arbitrary rotations $R_z(-\theta)$ are required. Because these rotations are in parallel, when the same Givens rotation is applied simultaneously many times (i.e. when changing bases in multiple dimensions) they can all be combined using Hamming weight phasing. For n orbitals, the circuit must be applied $n(n-1)/2$ times in each spin sector.

For diagonalizing the kinetic energy operator of a system with side length n in d dimensions, the discussion of [Appendix C.3.1](#) exactly carries over. The 1D (n -qubit) basis change circuit is applied several times in parallel, in d separate stages. Within each of the d stages, the 1D circuit is applied n^{d-1} times in parallel for the spinless system and $2n^{d-1}$ times in parallel for the spinful system. Hamming weight phasing can be applied across all these parallel circuits to reduce the number of arbitrary rotations.

D Costs per Trotter step

The number of T gates and arbitrary rotations per Trotter step figure significantly in the total costs of the two algorithms. For the fermionic swap network-based Trotter step, no T gates are directly required: instead, each application of the fermionic simulation gate requires up to 4 arbitrary rotations which must be compiled into T gates. On the other hand, for the split-operator algorithm, changing basis using either the FFT or Givens rotations may itself require T gates and arbitrary rotations. While the split-operator algorithm is in the momentum basis, evolution under the entire kinetic energy part of the Hamiltonian can be implemented by an arbitrary rotation on each qubit. While the algorithm is in the position basis, each interaction between spin-orbitals requires an arbitrary rotation, as well as an additional arbitrary rotation on each qubit to implement evolution under the single-qubit Z terms.

In this section, we discuss how to count the number of T gates and arbitrary rotations required per Trotter step for both the fermionic swap-based Trotter step and the split-operator Trotter step for all systems discussed in the paper.

D.1 The fermionic swap network

For the fermionic swap network simulating the uniform electron gas (jellium), the second-order (symmetric) Trotter step requires $4(N-1)^2$ arbitrary rotations for N qubits in the spinless case. Each application of the fermionic swap gate requires 4 arbitrary rotations (see [Figure 11](#)), and the fermionic swap gate is applied $2\binom{N}{2} = N(N-1)$ times, twice for each pair of orbitals because of the symmetric step. This would imply a cost at most $4N(N-1)$ arbitrary rotations. However, we can slightly reduce this cost by noting that (i) the arbitrary rotations in the first and final layers of the fermionic swap network can be merged in the symmetric step, and (ii) for the case with spin, many interactions do not occur between electrons in different spin sectors. If there are an equal number of interactions between same-spin and opposite-spin sectors at the start and end of the Trotter step, there are $3N-4$ arbitrary rotation gates total in the first and last layer of the step which may be merged.

The first of these observations allows us to eliminate $4(N-1)$ arbitrary rotations from the fermionic swap network Trotter step. This reduces the number of arbitrary rotations to $4(N-1)^2$ for the case of the uniform electron gas without spin. In the case of jellium with spin, the key for the second

observation is that the kinetic energy terms T_{pq} are zero if $\sigma_p \neq \sigma_q$, i.e. if p and q belong to different spin sectors. In this case, $\mathcal{F}_t(i_p, i_q)$ only requires 2 arbitrary rotations. Only considering the potential energy terms, the fermionic swap network requires $2N(N-1)$ arbitrary rotations. The kinetic energy only acts within the same spin sector. There are $\binom{N/2}{2} = N^2/8 - N/4$ pairs of orbitals within each spin sector requiring 2 arbitrary rotations for the kinetic energy terms, two spin sectors, and the symmetric Trotter step applies the corresponding rotations twice. In total then, the kinetic energy operator requires at most $N^2 - 2N$ arbitrary rotations. Combining this with the $2N^2 - 2N$ arbitrary rotations for the potential energy operator, the fermionic swap network requires at most $3N^2 - 4N$ arbitrary rotations for N qubits for the system with spin. Using the improvement of the first observation, there are at most $(N-1)(3N-4)$ arbitrary rotations to perform. Materials are nearly identical to jellium, but add up to N arbitrary rotations (one on each qubit) to the Trotter step.

For simulating the Hubbard model with spin, there are $N/2$ on-site interactions in the Hamiltonian (one between the two spin sectors for each of the $N/2$ spatial orbitals), and $2N$ hopping terms (two for each of the N spin-orbitals in the system). Each on-site interaction and each hopping term requires two arbitrary rotations to simulate. Symmetrizing for the second-order Trotter step adds an additional $4N$ arbitrary rotations for the hopping terms; however, the interaction terms in the symmetric step can be combined. Totalling these costs, at most $9N$ arbitrary rotations are required to simulate the Hubbard model.

D.2 The split-operator Trotter step

The costs for the split-operator Trotter step for the uniform electron gas and Hubbard model are as follows. For both systems, applying the kinetic energy operator in the momentum basis requires $2N$ arbitrary rotations, two for each qubit for the symmetric Trotter step. Because all terms in the potential energy operator commute, each rotation can be performed just once even for the symmetric step: $\binom{N}{2} + N$ arbitrary rotations are required at most. While the uniform electron gas saturates this requirement, the Hubbard model only needs $3N/2$ arbitrary rotations for the potential. As discussed in [Appendix C](#), because we work with a large number of Trotter steps, the cost of simulating terms in the symmetric split-operator Trotter step is approximately the sum of the cost of simulating the kinetic terms and the cost of simulating the potential terms. (For a single symmetric split-operator step, either the kinetic or the potential terms would be simulated twice, i.e. $\exp(-iHt) \approx \exp(-iTt/2)\exp(-iVt)\exp(-iTt/2)$ or $\exp(-iHt) \approx \exp(-iVt/2)\exp(-iTt)\exp(-iVt/2)$: when the number of steps is large, we can merge the halved term with the same term in the next Trotter step. We are thus neglecting in our counts the cost of the final simulation term, which cannot be merged in this way, but this is negligible when the number of Trotter steps is large.) In total, then, implementing the terms in the uniform electron gas Hamiltonian requires $\binom{N}{2} + 3N$ arbitrary rotations while the Hubbard model requires $7N/2$ arbitrary rotations.

We have not yet discussed the cost of changing bases. Changing between the position and momentum bases represents an additional cost for the split-operator Trotter step. Recall from [Appendix C](#) that the Fourier transform must be applied dM^{d-1} times to transform a spinless system of side length M in d dimensions, and $2dM^{d-1}$ times for the same system with spin. When the side length is a power of two, the FFT can be used, with a cost of 8 T gates for the 4-qubit case, 26 for the 8-qubit case, and 54 T gates and 4 arbitrary rotations for the 16-qubit case. Otherwise, the Givens rotation-based basis change algorithm must be used [29]. The basis change algorithm on M qubits requires $2\binom{M}{2}$ arbitrary rotations. In either case, the rotation from the position to the momentum basis and its reverse must be applied twice for the symmetric Trotter step.

Recall from [Appendix A](#) that Hamming weight phasing may be used to reduce the number of arbitrary rotations required for changing the basis, whether the FFT or Givens rotation algorithm is used. The basis must be applied twice in the symmetric split-operator Trotter step, i.e. each of the following costs will be doubled in our final accounting.

For the 16×16 FFT without spin, we can catalyze all non-Clifford rotations at a total cost of 2304 T and 64 Toffoli gates. The 16×16 FFT with spin uses 4608 T and 128 Toffoli gates.

To diagonalize the kinetic energy operator using Givens rotations with Hamming weight phasing on a d -dimensional grid of side length M , the costs are as follows. Without spin, $d \binom{M}{2} (\lfloor (d-1) \log_2 M \rfloor + 2)$ arbitrary rotations must be synthesized; with spin, $d \binom{M}{2} (\lfloor (d-1) \log_2 M \rfloor + 3)$ need to be synthesized. This introduces an additional cost in terms of T and Toffoli gates. Without spin, $d \binom{M}{2} (2M^{d-1} + 2 \lfloor (d-1) \log_2 M \rfloor)$ additional Toffoli gates are introduced by this procedure whereas with spin the number of additional Toffoli gates is $d \binom{M}{2} (4M^{d-1} + 2 \lfloor (d-1) \log_2 M \rfloor + 2)$.

D.3 Gate counts per Trotter step with Hamming weight phasing

The Trotter step gate counts we use are actually slightly smaller than those discussed in the previous two subsections. There are three reasons for this, which we describe in order of increasing complexity. First, we do not include zero-rotations. Second, for both the fermionic swap and split-operator Trotter steps, we neglect the cost of the final term in the simulation (the final layer, for the fermionic swap Trotter step), since those arbitrary rotations can be merged with the first term or layer of the next Trotter step. Third, we use Hamming weight phasing with a fixed number of ancilla qubits to combine equiangular rotations where possible. We discuss these aspects further in [Section 2](#).

E Improved second-order Trotter-Suzuki error bounds

Bounds on the error in Trotterized evolution typically assume that the evolution time t is small. Here, we present a tighter bound on the second-order Trotter error than that in Appendix B of [9], which, similar to the proof in that paper, does not require t to be small. Broadly, our proof is quite similar overall to that in [9], and we adopt similar notation so as to make it easier to identify where they differ. However, our result is tighter by constant factors, matching the factor of $1/12$ from the Baker-Campbell-Hausdorff expansion discussed in [12]. Additionally, we are able to keep the norm on the outside of some of the sums, which is another reason it is tighter than the bound in [9]. It has been pointed out that the bound described here can in principle be reached from the work in [46] using slightly different methods.

The second-order Trotter formula for e^{-iHt} is given by Eq. (5),

$$U = e^{-iHt} \approx \left(\prod_{\ell=1}^L e^{-iH_\ell t/2} \prod_{\ell=L}^1 e^{-iH_\ell t/2} \right) \equiv U^{\text{TS}} = e^{-iH_{\text{eff}}t} \quad (37)$$

for the Hamiltonian $H = \sum_{\ell=1}^L H_\ell$. We first consider the error when H is a sum of only two terms. We write $Ht = A + B$, and let $H(x) = B + (1-x)A$. Then the error in the Trotter formula for A and B can be written as

$$\begin{aligned} & \left\| \exp(-iA/2) \exp(-iB) \exp(-iA/2) - \exp(-iH(0)) \right\| \\ &= \left\| \int_0^1 \partial_x (\exp(-ixA/2) \exp(-iH(x)) \exp(-ixA/2)) dx \right\|. \end{aligned} \quad (38)$$

Now, define $A(u, x) = \exp(iH(x)u)A \exp(-iH(x)u)$. Then

$$\begin{aligned} & \partial_x [\exp(-ixA/2) \exp(-iH(x)) \exp(-ixA/2)] \\ &= \exp(-ixA/2) \left(-\frac{i}{2} A \exp(-iH(x)) + \partial_x \exp(-iH(x)) - \frac{i}{2} \exp(-iH(x)) A \right) \exp(-ixA/2). \end{aligned} \quad (39)$$

In bounding this, [9] use

$$\begin{aligned} \partial_x \exp(-iH(x)) &= \int_0^1 \exp(-i(1-u)H(x)) (iA) \exp(-iuH(x)) du \\ &= i \exp(-iH(x)) \int_0^1 A(u, x) du. \end{aligned} \quad (40)$$

Additionally, $A \exp(-iH(x)) = \exp(-iH(x))A(1, x)$. That then gives the expression for the derivative

$$\begin{aligned} & \partial_x (\exp(-ixA/2) \exp(-iH(x)) \exp(-ixA/2)) \\ &= \frac{i}{2} \exp(-ixA/2) \exp(-iH(x)) \left(-A - A(1, x) + 2 \int_0^1 A(u, x) du \right) \exp(-ixA/2). \end{aligned} \quad (41)$$

Next, using $A(0, x) = A$ and

$$A(u, x) = A(0, x) + uA'(0, x) + \int_0^u (u-s)A''(s, x) ds, \quad (42)$$

where A' and A'' are the respectively the first and second derivatives of $A(u, x)$ with respect to u , we have that the terms in parentheses in Eq. (41) are equivalent to

$$\begin{aligned} -A - A(1, x) + 2 \int_0^1 A(u, x) du &= -A - A - A'(0, x) - \int_0^1 (1-s)A''(s, x) ds \\ &+ 2 \int_0^1 A du + 2 \int_0^1 uA'(0, x) du + 2 \int_0^1 du \int_0^u ds (u-s)A''(s, x) \\ &= - \int_0^1 (1-s)A''(s, x) ds + 2 \int_0^1 ds \int_s^1 du (u-s)A''(s, x) \\ &= - \int_0^1 (1-s)A''(s, x) ds + \int_0^1 (1-s^2 - 2s + 2s^2)A''(s, x) ds \\ &= \int_0^1 (s^2 - s)A''(s, x) ds \end{aligned} \quad (43)$$

Evaluating the derivatives gives

$$A'(s, x) = i(H(x)A(s, x) - A(s, x)H(x)) \quad (44)$$

and

$$\begin{aligned} A''(s, x) &= -((H(x))^2 A(s, x) + A(s, x)(H(x))^2 - 2H(x)A(s, x)H(x)) \\ &= -\exp(iH(x)s) ((H(x))^2 A + A(H(x))^2 - 2H(x)AH(x)) \exp(-iH(x)s). \end{aligned} \quad (45)$$

Substituting in the expression for $H(x)$ gives

$$\begin{aligned} (H(x))^2 A + A(H(x))^2 - 2H(x)AH(x) &= (B^2 + (1-x)(AB + BA) + (1-x)^2 A^2)A \\ &+ A(B^2 + (1-x)(AB + BA) + (1-x)^2 A^2) - 2(BAB + (1-x)(BAA + AAB) + (1-x)^2 A^3) \\ &= B^2 A + AB^2 - 2BAB + (1-x)(ABA + BAA + AAB + ABA - 2BAA - 2AAB) \\ &= B^2 A + AB^2 - 2BAB + (1-x)(2ABA - BAA - AAB) \\ &= [[A, B], B] + (1-x)[[A, B], A]. \end{aligned} \quad (46)$$

Combining these results in Eq. (41), we therefore have

$$\begin{aligned} & \partial_x (\exp(-ixA/2) \exp(-iH(x)) \exp(-ixA/2)) \\ &= \frac{i}{2} \exp(-ixA/2) \left(\int_0^1 ds (s - s^2) \exp(iH(x)(s-1)) ([A, B], B) + (1-x)[[A, B], A] \right) \exp(-ixA/2) \times \\ & \quad \exp(-ixA/2). \end{aligned} \quad (47)$$

Therefore,

$$\begin{aligned} \|\partial_x (\exp(-ixA/2) \exp(-iH(x)) \exp(-ixA/2))\| &\leq \frac{1}{2} \int_0^1 ds (s - s^2) \|[[A, B], B] + (1-x)[[A, B], A]\| \\ &= \frac{1}{12} \|[[A, B], B] + (1-x)[[A, B], A]\|. \end{aligned} \quad (48)$$

This has given a tighter bound by a factor of 1/12 compared to that given in [9], which matches the factor from a BCH expansion. Lastly, integrating over x gives

$$\begin{aligned} &\|\exp(-iA/2) \exp(-iB) \exp(-iA/2) - \exp(-iH(0))\| \\ &\leq \int_0^1 \|\partial_x (\exp(-ixA/2) \exp(-iH(x)) \exp(-ixA/2))\| dx \\ &\leq \frac{1}{12} \int_0^1 \|[[A, B], B] + (1-x)[[A, B], A]\| dx \\ &\leq \frac{1}{12} \|[[A, B], B]\| + \frac{1}{24} \|[[A, B], A]\|. \end{aligned} \quad (49)$$

Next, one can use the iterative approach from [9]. Let

$$U_j = \exp\left(-i \sum_{k=j}^L H_k t\right) \quad (50)$$

$$U_j^{TS} = \exp\left(-iH_j \frac{t}{2}\right) \dots \exp\left(-iH_{L-1} \frac{t}{2}\right) \exp\left(-iH_L \frac{t}{2}\right) \exp\left(-iH_{L-1} \frac{t}{2}\right) \dots \exp\left(-iH_j \frac{t}{2}\right). \quad (51)$$

Using the triangle inequality

$$\begin{aligned} &\|U_{j-1} - U_{j-1}^{TS}\| \\ &\leq \left\| U_{j-1} - \exp\left(-iH_{j-1} \frac{t}{2}\right) U_j \exp\left(-iH_{j-1} \frac{t}{2}\right) \right\| + \left\| \exp\left(-iH_{j-1} \frac{t}{2}\right) U_j \exp\left(-iH_{j-1} \frac{t}{2}\right) - U_{j-1}^{TS} \right\| \\ &= \left\| U_{j-1} - \exp\left(-iH_{j-1} \frac{t}{2}\right) U_j \exp\left(-iH_{j-1} \frac{t}{2}\right) \right\| + \|U_j - U_j^{TS}\|. \end{aligned} \quad (52)$$

Noting that $U_1 = U$ and $U_1^{TS} = U^{TS}$, and $U_L = U_L^{TS}$, we have

$$\|U - U^{TS}\| \leq \sum_{b=1}^{L-1} \left\| U_b - \exp\left(-iH_b \frac{t}{2}\right) U_{b+1} \exp\left(-iH_b \frac{t}{2}\right) \right\|. \quad (53)$$

Now, using the bounds above with

$$A = H_b t, \quad B = \sum_{c>b} H_c t, \quad (54)$$

we have

$$[[A, B], B] = t^3 \sum_{c>b} \sum_{a>b} [[H_b, H_c], H_a], \quad [[A, B], A] = t^3 \sum_{c>b} [[H_b, H_c], H_b]. \quad (55)$$

This gives

$$\|U - U^{TS}\| \leq \frac{t^3}{12} \sum_{b=1}^{L-1} \left(\left\| \sum_{c>b} \sum_{a>b} [[H_b, H_c], H_a] \right\| + \frac{1}{2} \left\| \sum_{c>b} [[H_b, H_c], H_b] \right\| \right). \quad (56)$$

Lastly, since we are using the evolution to measure the energy, we need to bound that in terms of the difference in the unitaries. The difference in the eigenvalues of unitaries may be upper bounded as [88]

$$|e^{-iEt} - e^{-iE_{TS}t}| \leq \|U - U^{TS}\|. \quad (57)$$

Using $\Delta \equiv \|U - U^{TS}\|$, for $\Delta^2 \leq 2$

$$|E - E_{TS}|t \leq \arctan\left(\Delta \frac{\sqrt{4 - \Delta^2}}{2 - \Delta^2}\right) = \Delta + \frac{\Delta^3}{24} + \mathcal{O}(\Delta^5). \quad (58)$$

F Phase estimation circuit primitives from a Trotter step

Phase estimation typically is assumed to use a controlled unitary rotation, however, [43] showed that phase estimation can be made more efficient by eschewing controlled quantum evolutions in exchange for circuits that control the direction of the time evolution. That is to say, we want a circuit such that $|c\rangle|\psi\rangle \mapsto |c\rangle e^{-iHt(-1)^c}|\psi\rangle$. We build such a circuit by exploiting the structure of the Trotter decomposition. If the symmetric Trotter formula is used then we can simulate the inverse time evolution by flipping the sign of each evolution in the decomposition. That is to say, if $U_1 := \prod_{j=1}^M e^{-iH_j t} \prod_{j=M}^1 e^{-iH_j t}$ then $U_1^\dagger = \prod_{j=1}^M e^{iH_j t} \prod_{j=M}^1 e^{iH_j t}$. This does not hold for non-symmetric splittings such as the lowest-order Trotter formula. Thus, we simply need to find a circuit for controllably reversing the direction of each H_j . Such circuits are simple to produce at no extra cost (compared to the Trotter step circuits) for the potential terms; this allows implementation of the directionally controlled circuit for the split-operator Trotter step. For that case, the directionally-controlled circuits can be constructed by adding a controlled-NOT gate (controlled on the phase estimation ancilla, targeting the qubit in question) on either side of each of the arbitrary-angle R_z gates in the circuit. We show the directionally-controlled version of \mathcal{V}_t , the circuit for simulating the interaction terms $-V_{pq}Z_p Z_q/4$, in Figure 10.

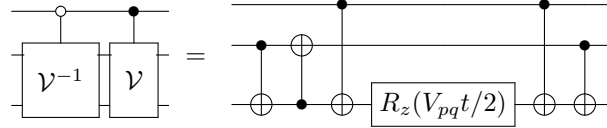


Figure 10: Directionally controlled fault-tolerant circuit for interaction terms.

The number of arbitrary rotations and T gates for the directionally-controlled circuit is thus the same as for the bare Trotter step. However, integrating the kinetic energy terms between electrons in a constant spin manifold into this, and hence constructing the directionally-controlled fermionic swap Trotter step, is more challenging.

The simulation circuit in Figure 11 achieves this directional control for the fermionic swap network. It follows from the circuit for $\mathcal{F}_t(i_p, i_q)$ after noting that conjugating R_z gates with NOT gates flips the direction of rotation. The result given in the circuit follows from this observation and elementary identities for simplifying chains of CNOT gates. Importantly, the number of arbitrary rotations required for the circuit is identical to that of the corresponding uncontrolled circuit.

G Trotter error numerics

Thus far, we have discussed in the appendices the costs of single Trotter steps. We have presented this cost in terms of arbitrary rotations and T gates. In this appendix, we discuss our computational

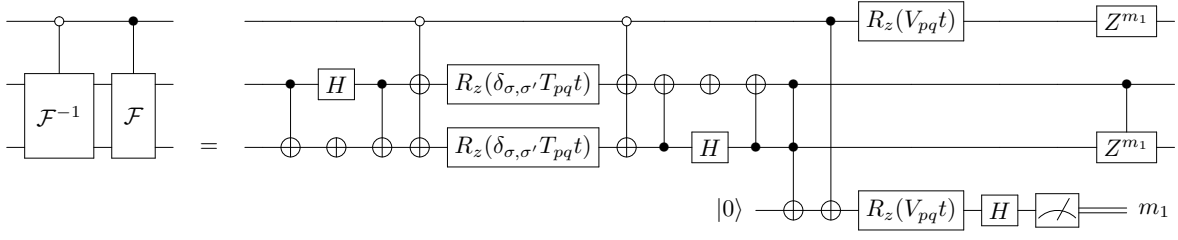


Figure 11: How to implement simultaneous \mathcal{F} operations followed by fermionic swaps. Note that the phasing operation on the control has negligible cost. Although it may appear that there are $\mathcal{O}(M_k \cdot n^2)$ phasing operations applied to the control in the course of evaluating a Trotter step, all of the control phasing operations can be merged into a single phasing operation.

strategy for determining the errors in Trotterizing evolution, enabling us to compute the total cost of the two algorithms using phase estimation. We do this by computing the second-order Trotter error norm W as in Eq. (6),

$$\|e^{-iHt} - e^{-iH_{\text{eff}}t}\| \leq \frac{t^3}{12} \sum_{b=1}^{L-1} \left(\left\| \sum_{c>b} \sum_{a>b} [[H_b, H_c], H_a] \right\| + \frac{1}{2} \left\| \sum_{c>b} [[H_b, H_c], H_b] \right\| \right) \equiv Wt^3, \quad (59)$$

and thereby bounding the shifts in the energy as in Eq. (6). Recall that we call W the Trotter error norm. The key difficulty in bounding the Trotter error is the sheer number of terms in the Hamiltonian: if there are L terms in the Hamiltonian, then W is a sum of L^3 double commutators. Given that the number of terms in the uniform electron gas Hamiltonian scales quadratically with the number of spin-orbitals N , this represents at most $\mathcal{O}(N^6)$ double commutators that must be evaluated. In fact, the number of double commutators that are non-zero in this expansion is significantly smaller than this as each term in the nested commutators must share at least one index for the commutator to be non-zero. This results in $\mathcal{O}(N^4)$ non-trivial terms in the expansion of the error. By comparison, the locality of the Hubbard model means there are only $\mathcal{O}(N)$ non-zero double commutators.

The Trotter error norms are computed as follows. For each system, the Trotterization schemes of Section 2 (and further explained in Appendix C and Appendix B) give the ordered decomposition $\sum_a H_a$ in which the terms in the Hamiltonian are simulated. Note that the number of terms in the Trotter order is only the same as the number of terms in the Hamiltonian up to constant factors. For the fermionic swap network, the terms $a_i^\dagger a_j$ and $a_j^\dagger a_i$ can be grouped together at each stage, as can $a_i^\dagger a_j^\dagger a_i a_j$, $a_i^\dagger a_i$, and $a_j^\dagger a_j$. For the split-operator algorithm, all the terms in the kinetic energy can be grouped together in the sum, as can all the terms in the potential energy operator. As a result, the Trotter error norm for the split-operator Trotter step has only two terms. When the kinetic energy operator is simulated before the potential, the second-order Trotter error norm for the split-operator step is

$$W_{T+V} = \|[V, [V, T]]\| + \frac{1}{2} \|[T, [V, T]]\|. \quad (60)$$

On the other hand, when the potential is simulated first, the resulting Trotter error norm is

$$W_{V+T} = \|[T, [V, T]]\| + \frac{1}{2} \|[V, [V, T]]\|. \quad (61)$$

Depending on the norms of T and V this can yield different bounds on the Trotter error. So long as the number of Trotter steps r is large, the number of gates in the simulation does not depend significantly on whether T or V is simulated first, and the user can freely choose the case yielding the smaller error. The fermionic swap network is less flexible in this regard.

Given these ordered sums, we then compute the double commutators, and total them to find the Trotter error norm. The order in the sums is specified by the Trotter step. Computing the error norms

UEG System	System Logical Qubits	W_{FS}	W_{SO}
4 × 4 spinless	16	2.2e+00	9.9e-01
5 × 5 spinless	25	6.5e+00	2.3e+00
6 × 6 spinless	36	3.2e+01	1.3e+01
7 × 7 spinless	49	5.7e+01	1.9e+01
8 × 8 spinless	64	1.9e+02	7.0e+01
9 × 9 spinless	81	2.8e+02	8.5e+01
10 × 10 spinless	100	6.9e+02	2.5e+02
11 × 11 spinless	121	9.8e+02	2.6e+02
12 × 12 spinless	144	2.0e+03	6.7e+02
13 × 13 spinless	169	2.6e+03	6.6e+02
14 × 14 spinless	196	4.9e+03	1.6e+03
15 × 15 spinless	225	6.2e+03	1.4e+03
16 × 16 spinless	256	1.0e+04	3.2e+03
4 × 4 spinful	32	1.4e+01	6.9e+00
5 × 5 spinful	50	3.3e+01	1.4e+01
6 × 6 spinful	72	1.6e+02	6.9e+01
7 × 7 spinful	98	2.6e+02	9.8e+01
8 × 8 spinful	128	8.5e+02	3.4e+02
9 × 9 spinful	162	1.2e+03	4.0e+02
10 × 10 spinful	200	3.0e+03	1.1e+03
11 × 11 spinful	242	3.8e+03	1.2e+03
12 × 12 spinful	288	8.3e+03	3.0e+03

Table 4: Trotter error data for simulating the uniform electron gas, for both the fermionic swap (FS) and split-operator (SO) Trotter steps, at Wigner-Seitz radius $r_s = 10$ in 2D. W_X denotes the Trotter error norm for algorithm X . Hartree-Fock energies E_{HF} serve as an upper bound on the true ground state energies for the systems considered. Energies are in units of Hartree (Ha), Trotter error norms are in cubed Hartree (Ha^3).

for larger sizes requires significant computational resources. For example, the Hamiltonian for one of the larger system sizes we study for the uniform electron gas (12×12 with spin) has 47,952 terms in the Hamiltonian and 44,496 terms in the Trotter ordering for the fermionic swap network.

This means that computing W requires the evaluation of as many as 10^{14} double commutators. We partially ameliorate this by performing preprocessing to determine which double commutators $[[H_b, H_c], H_a]$ are trivially zero and do not need to be evaluated. The Trotter error norm for the larger systems can require significant memory to load into memory (> 100 GB). We work around this issue by parallelizing computation of across orthogonal components as follows. Writing the error norm as a sum of Pauli strings $W = \sum_i a_i \|P_i\|$, let the smallest qubit index P_i acts on be q_i . Then we can compute the sum of $|a_i|$ for which q_i is some fixed value (i.e., the triangle inequality bound for a given qubit index), and the sum over q_i of these values is equal to the full Trotter error norm. This allows us to control the maximum memory requirement of any partial computation.

Trotter error norms in units of cubed Hartree (Ha^3) are included for the uniform electron gas in 2D in Table 4 and in 3D in Table 5. We additionally plot the Trotter error norm for these systems in Figure 12. Table 6 shows Trotter error norm data for the Hubbard model in 2D (with spin). The same data is plotted in Figure 13 with a power law fit. We include the same data for a range of Wigner-Seitz radii in 3D in Table 7 and for a series of materials in Table 8. For the materials, we include in the table the Hartree-Fock energies computed in PySCF [89] which we use for our relative precision energy target.

All numerics (for all systems and for the different Trotter steps) were performed using code which we have contributed to the open-source package OpenFermion [77].

UEG System	System Logical Qubits	W_{FS}	W_{SO}
$3 \times 3 \times 3$ spinless	27	7.1e-03	2.4e-03
$4 \times 4 \times 4$ spinless	64	1.7e-01	6.8e-02
$5 \times 5 \times 5$ spinless	125	4.7e-01	1.6e-01
$6 \times 6 \times 6$ spinless	216	2.6e+00	9.7e-01
$2 \times 2 \times 2$ spinful	16	2.4e-03	1.5e-03
$3 \times 3 \times 3$ spinful	54	2.6e-02	1.3e-02
$4 \times 4 \times 4$ spinful	128	6.0e-01	2.5e-01
$5 \times 5 \times 5$ spinful	250	1.5e+00	5.9e-01

Table 5: Uniform electron gas (UEG) Trotter error data, for both the fermionic swap (FS) and split-operator (SO) Trotter steps, at Wigner-Seitz radius $r_s = 10$ in 3D with varying system side length. W_X denotes the Trotter error norm for algorithm X . Hartree-Fock energies E_{HF} serve as an upper bound on the true ground state energies for the systems considered. Energies are in units of Hartree (Ha), Trotter error norms are in cubed Hartree (Ha^3).

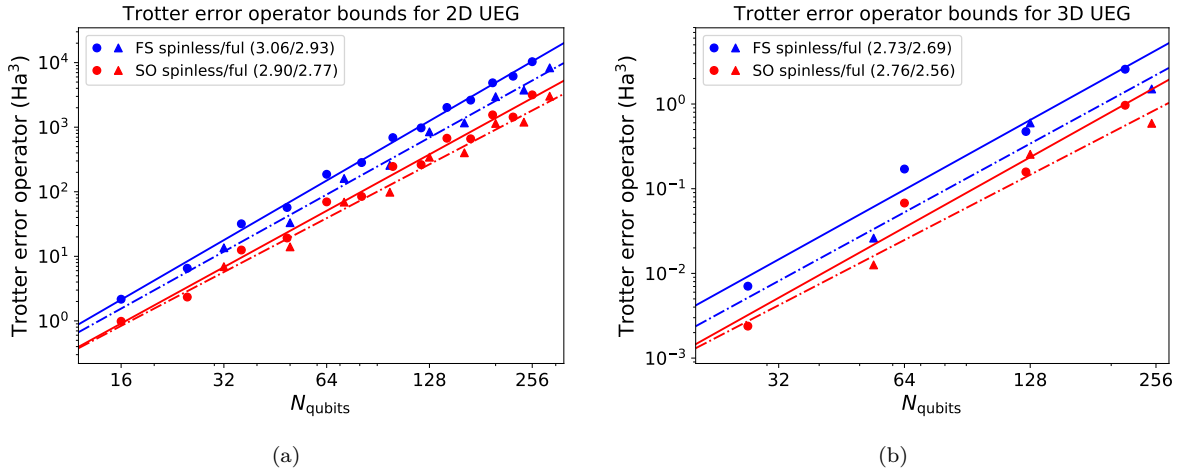


Figure 12: Trotter error norms for the uniform electron gas (UEG) Hamiltonian in (a) 2D and (b) 3D, for both spinful and spinless systems, in units of cubed Hartree (Ha^3). Lines show power law fits to the data, with the exponents of the fits shown in the caption for each system.

H T-count minimization

Continuing from the discussion in Section 3, we present in this appendix our final numerics to compute the total number of T and Toffoli gates required for our simulations. We wish to minimize the total number of T and Toffoli gates needed to determine the ground state energy to precision $\Delta E \geq \Delta E_{TS} + \Delta E_{PE} + \Delta E_{HT}$ using phase estimation. ΔE_{TS} accounts for Trotter-Suzuki errors, ΔE_{PE} for phase estimation errors, and ΔE_{HT} for circuit synthesis errors. Throughout, we convert Toffoli to T gates at a cost of 2 T gates per Toffoli gate [81].

The tradeoff between these error sources is non-trivial because the T-costs differ significantly as a function of the desired accuracy for each type of error. We discuss how to compute the T-cost as a function of these errors, and give the results of its numerical minimization subject to Eq. (11) for a variety of system sizes, in Section 3.2. Our full objective function there was given in Eq. (12), as

$$(N_r N_{HT} + N_d) N_{PE} \approx \frac{0.76\pi\sqrt{W}}{\Delta E_{PE}\sqrt{\Delta E_{TS}}} \left(N_r \left[1.15 \log_2 \left(\frac{N_r\sqrt{W}}{\Delta E_{HT}\sqrt{\Delta E_{TS}}} \right) + 9.2 \right] + N_d \right),$$

Hubbard System	System Log. Qubits	$W_{FS}, U = 4$	$W_{SO}, U = 4$	$W_{FS}, U = 8$	$W_{SO}, U = 8$
4×4	32	1.9e+03	1.4e+03	4.3e+03	4.1e+03
5×5	50	2.9e+03	2.1e+03	6.7e+03	6.4e+03
6×6	72	4.2e+03	3.1e+03	9.6e+03	9.2e+03
7×7	98	5.7e+03	4.2e+03	1.3e+04	1.3e+04
8×8	128	7.5e+03	5.5e+03	1.7e+04	1.6e+04
9×9	162	9.4e+03	6.9e+03	2.2e+04	2.1e+04
10×10	200	1.2e+04	8.5e+03	2.7e+04	2.6e+04
11×11	242	1.4e+04	1.0e+04	3.2e+04	3.1e+04
12×12	288	1.7e+04	1.2e+04	3.8e+04	3.7e+04
13×13	338	2.0e+04	1.4e+04	4.5e+04	4.3e+04
14×14	392	2.3e+04	1.7e+04	5.2e+04	5.0e+04
15×15	450	2.6e+04	1.9e+04	6.0e+04	5.8e+04
16×16	512	3.0e+04	2.2e+04	6.8e+04	6.6e+04

Table 6: Hubbard model Trotter error data for $t = 1$, with $U = 4$ and $U = 8$, for both the fermionic swap (FS) and split-operator (SO) Trotter steps, for a range of system sizes. W_X denotes the Trotter error norm for algorithm X .

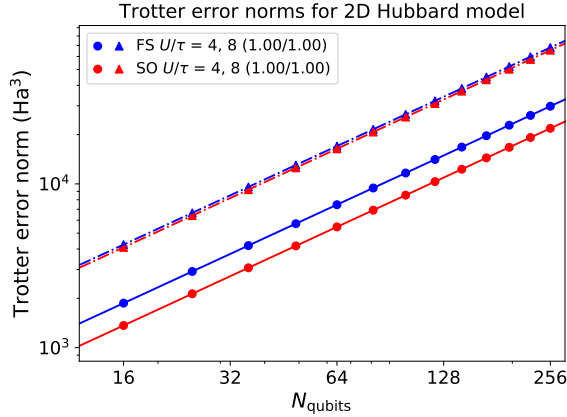


Figure 13: Trotter error norms for the Hubbard model in cubed Hartree (Ha^3). Lines show power law fits to the data, with the exponents of the fits shown in the caption for each system.

where W is the Trotter error norm, N_r is the number of arbitrary rotations per Trotter step, and N_d is the number of “direct” T and Toffoli gates (those T/Toffoli gates which appear in the circuit before synthesis) multiplied by N_{PE} .

Data on the Trotter error norms W for all systems and algorithms are given in Appendix G, and our approach to computing the numbers of arbitrary rotations N_r and T/Toffoli gates which appear in the circuit before compiling arbitrary rotations (“direct” T/Toffoli gates) N_d is discussed in Section 2 as well as in Appendix D. We work with two total precision targets ΔE for all the systems we analyze. In the first, we consider a relative precision, with the total precision ΔE set as a fraction (0.5%) of the ground state energy. This value is comparable to errors associated with the sign problem in quantum Monte Carlo [82, 83]. The precision ΔE in this case scales with the system size. In the second case, we consider an absolute precision target set depending on the type of system, but independently of the particular system size.

For the relative precision case for the jellium systems, since we cannot directly compute the ground state energy, we take ΔE to be 0.5% of the system energy, computed as the number of electrons multiplied by the energy per electron using the correlation energy estimates of [90]. In atomic units,

the energy per electron takes the form

$$\frac{\tilde{E}_0(r_s)}{\eta} \approx \frac{3}{5} \left(\frac{9\pi}{4} \right)^{2/3} \frac{1}{r_s^2} - \frac{3}{2} \left(\frac{9\pi}{4} \right)^{1/3} \frac{1}{r_s} + a \ln \left(1 + \frac{b}{r_s} + \frac{b}{r_s^2} \right). \quad (62)$$

We use the values $a = (\ln 2 - 1)/(2\pi^2)$ and $b = 20.4562557$ over the entire range of densities. This closely agrees with quantum Monte Carlo simulations in the low-density regime, and lower-bounds the magnitude of the ferromagnetic energy in the high-density regime. We then take ΔE to be 0.5% of this value, after multiplying by the number of electrons (η in the expression above). We use these energies in computing the T-cost for the uniform electron gas at a Wigner-Seitz radius of 10 Bohr radii in Figure 1, as well as for a range of Wigner-Seitz radii below in Figure 14. For the material systems, we take as our relative precision target the Hartree-Fock energy E_{HF} (given in Appendix G). We compute these using PySCF [89]. The true ground state energy is $E_0 \leq \tilde{E}_0 = E_{\text{HF}}$ in general, and $E_{\text{HF}} < 0$ for all material systems considered, so choosing $\Delta E = 0.005 E_{\text{HF}}$ ensures that the precision better than 0.5% of the magnitude of the ground state energy.

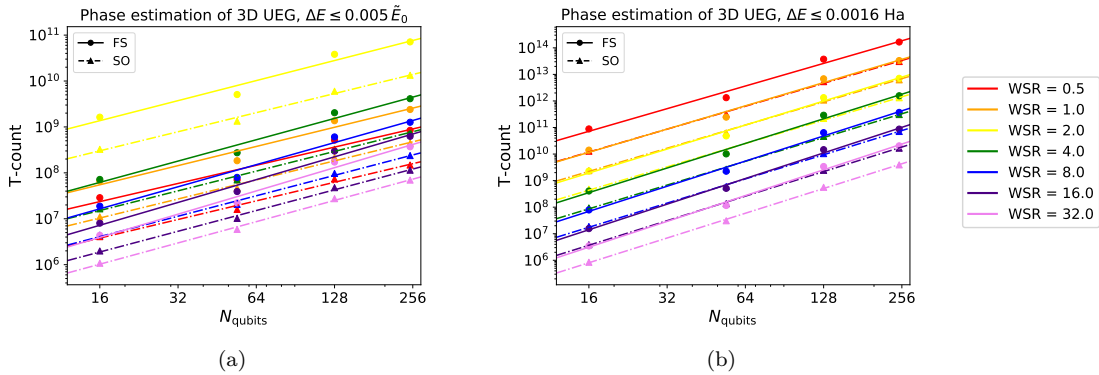


Figure 14: Loose but rigorous upper bound on the number of T gates in the circuit which performs our Trotterized phase estimation simulation of the 3D uniform electron gas, for a range of values of the Wigner-Seitz radius, to (a) within half a percent of the total system energy and to (b) absolute precision $\Delta E = 0.0016$ Hartree. The colors of the rainbow (red through violet) correspond to logarithmically spaced Wigner-Seitz radii from 0.5 to 32 Bohr radii (0.5, 1, 2, ..., 16, 32). Lines show power law fits to the data. 14 ancilla qubits are used for Hamming weight phasing for all systems.

Finally, for our relative precision target for the Hubbard model, we choose 1.02 as an upper bound on the magnitude of the energy per site at $U/\tau = 4$ and 0.74 as an upper bound on the magnitude of the energy per site at $U/\tau = 8$, following Table V of [39]; again, this ensures that the precision tolerance is smaller than 0.5% of the total ground state energy. For our absolute precision targets, we take $\Delta E = 0.0016$ Hartree for jellium and the material systems (this precision is known as “chemical accuracy”), and $\Delta E = \tau/100$ for the Hubbard model.

A possible criticism of the results of the optimization for the Hubbard model is that the number of Trotter steps decreases with system size in the relative precision case. Because of this, the scalings we present in Figure 2a do not hold in the asymptotic limit; rather, they only apply for a finite range of system sizes. We perform a power-law fit on the number of Trotter steps for the Hubbard model. This indicates that the system size at which the number of Trotter steps required goes to one is hundreds of thousands of spin-orbitals, even in the smallest case. We plot the number of Trotter steps with the power-law fit for the Hubbard model in Figure 15.

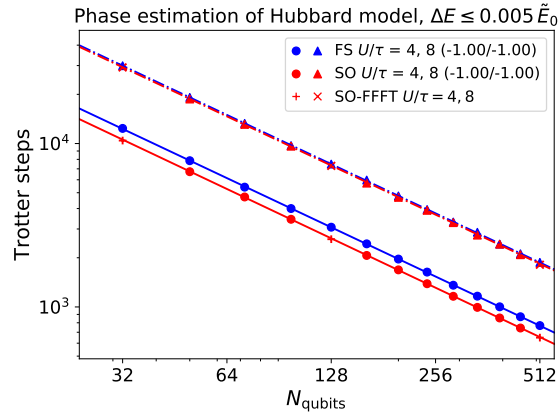


Figure 15: Power-law fit to the number of Trotter steps used for phase estimation, with 14 ancilla qubits used for Hamming weight phasing, for the Hubbard model in 2D. The exponents of the fit are included in the caption. The fits intercept 1 well above 100,000 system qubits, far above the range of system sizes we consider.

UEG System	System Logical Qubits	Wigner-Seitz Radius	W_{FS}	W_{SO}
$2 \times 2 \times 2$ spinful	16	0.125	1.0e+07	9.2e+05
$2 \times 2 \times 2$ spinful	16	0.25	2.1e+05	3.1e+04
$2 \times 2 \times 2$ spinful	16	0.5	5.1e+03	1.1e+03
$2 \times 2 \times 2$ spinful	16	1.0	1.4e+02	4.5e+01
$2 \times 2 \times 2$ spinful	16	2.0	4.4e+00	2.0e+00
$2 \times 2 \times 2$ spinful	16	4.0	1.6e-01	8.7e-02
$2 \times 2 \times 2$ spinful	16	8.0	6.5e-03	3.9e-03
$2 \times 2 \times 2$ spinful	16	16.0	3.1e-04	1.9e-04
$2 \times 2 \times 2$ spinful	16	32.0	1.6e-05	1.1e-05
$3 \times 3 \times 3$ spinful	54	0.125	1.9e+07	3.4e+06
$3 \times 3 \times 3$ spinful	54	0.25	5.1e+05	1.4e+05
$3 \times 3 \times 3$ spinful	54	0.5	1.6e+04	6.0e+03
$3 \times 3 \times 3$ spinful	54	1.0	6.0e+02	2.6e+02
$3 \times 3 \times 3$ spinful	54	2.0	2.5e+01	1.2e+01
$3 \times 3 \times 3$ spinful	54	4.0	1.2e+00	5.8e-01
$3 \times 3 \times 3$ spinful	54	8.0	6.6e-02	3.2e-02
$3 \times 3 \times 3$ spinful	54	16.0	3.8e-03	1.8e-03
$3 \times 3 \times 3$ spinful	54	32.0	2.3e-04	1.1e-04
$4 \times 4 \times 4$ spinful	128	0.125	4.1e+08	6.8e+07
$4 \times 4 \times 4$ spinful	128	0.25	1.1e+07	2.7e+06
$4 \times 4 \times 4$ spinful	128	0.5	3.6e+05	1.2e+05
$4 \times 4 \times 4$ spinful	128	1.0	1.3e+04	5.2e+03
$4 \times 4 \times 4$ spinful	128	2.0	5.7e+02	2.3e+02
$4 \times 4 \times 4$ spinful	128	4.0	2.8e+01	1.2e+01
$4 \times 4 \times 4$ spinful	128	8.0	1.5e+00	6.4e-01
$4 \times 4 \times 4$ spinful	128	16.0	8.7e-02	3.7e-02
$4 \times 4 \times 4$ spinful	128	32.0	5.2e-03	2.2e-03
$5 \times 5 \times 5$ spinful	250	0.125	4.7e+08	1.0e+08
$5 \times 5 \times 5$ spinful	250	0.25	1.5e+07	4.6e+06
$5 \times 5 \times 5$ spinful	250	0.5	5.7e+05	2.0e+05
$5 \times 5 \times 5$ spinful	250	1.0	2.4e+04	8.9e+03
$5 \times 5 \times 5$ spinful	250	2.0	1.2e+03	4.5e+02
$5 \times 5 \times 5$ spinful	250	4.0	6.5e+01	2.5e+01
$5 \times 5 \times 5$ spinful	250	8.0	3.8e+00	1.5e+00
$5 \times 5 \times 5$ spinful	250	16.0	2.3e-01	8.8e-02
$5 \times 5 \times 5$ spinful	250	32.0	1.4e-02	5.4e-03

Table 7: Uniform electron gas (UEG) Trotter error data, for both the fermionic swap (FS) and split-operator (SO) Trotter steps, for a range of systems in 3D with varying Wigner-Seitz radius and side length. W_X denotes the Trotter error norm for algorithm X . Hartree-Fock energies E_{HF} serve as an upper bound on the true ground state energies for the systems considered. Energies are in units of Hartree (Ha), Trotter error norms are in cubed Hartree (Ha^3). For those systems for which E_{HF} is positive, it is either exact or sufficiently large such that it is still a good proxy for E_0 for the purposes of relative precision.

Material System	System Logical Qubits	W_{FS}	W_{SO}	E_{HF}
Li $2 \times 2 \times 2$	16	1.3e+02	4.1e+01	-3.0e+00
Li $3 \times 3 \times 3$	54	4.7e+03	2.4e+03	-4.7e+00
Li $4 \times 4 \times 4$	128	3.3e+05	1.4e+05	-5.1e+00
Diamond $2 \times 2 \times 2$	16	9.7e-01	5.4e-01	-1.5e+01
Diamond $3 \times 3 \times 3$	54	3.0e+02	1.5e+02	-2.5e+01
Diamond $4 \times 4 \times 4$	128	6.9e+03	3.1e+03	-2.8e+01
Si $2 \times 2 \times 2$	16	3.1e-01	1.6e-01	-5.2e+01
Si $3 \times 3 \times 3$	54	6.6e+01	3.3e+01	-7.7e+01
Si $4 \times 4 \times 4$	128	1.2e+03	5.4e+02	-8.8e+01
LiH $2 \times 2 \times 2$	16	4.0e-01	2.4e-01	-2.6e+00
LiH $3 \times 3 \times 3$	54	1.4e+02	7.1e+01	-3.8e+00
LiH $4 \times 4 \times 4$	128	3.5e+03	1.6e+03	-4.4e+00
Graphite $2 \times 2 \times 2$	16	1.0e+02	1.4e+01	-2.8e+01
Graphite $3 \times 3 \times 3$	54	3.5e+03	4.1e+02	-3.4e+01
Graphite $4 \times 4 \times 4$	128	1.9e+05	2.9e+04	-5.2e+01

Table 8: Trotter error data for both the fermionic swap (FS) and split-operator (SO) Trotter steps for a range of materials in 3D, including Hartree-Fock energies. W_X denotes the Trotter error norm for algorithm X . Energies are in units of Hartree (Ha), Trotter error norms are in cubed Hartree (Ha^3).