

Phase Estimation with Partially Randomized Time Evolution

Jakob Günther^{1,3}, Freek Witteveen^{1,2,*}, Alexander Schmidhuber³, Marek Miller¹,
Matthias Christandl¹ and Aram W. Harrow³

¹*Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5,
2100 Copenhagen, Denmark*

²*QuSoft and CWI, Science Park 123, 1098 XG Amsterdam, The Netherlands*

³*Center for Theoretical Physics - a Leinweber Institute, MIT, Cambridge, Massachusetts 02139, USA*



(Received 12 May 2025; revised 16 September 2025; accepted 31 March 2026; published 19 May 2026)

Quantum phase estimation combined with Hamiltonian simulation is the most promising algorithmic framework to computing ground-state energies on quantum computers. Its main computational overhead derives from the Hamiltonian simulation subroutine. In this paper we use randomization to speed up product formulas, one of the standard approaches to Hamiltonian simulation. We propose partially randomized Hamiltonian simulation methods in which some terms are kept deterministically and others are randomly sampled. We perform a detailed resource estimate for single-ancilla phase estimation using partially randomized product formulas for benchmark systems in quantum chemistry and obtain orders-of-magnitude improvements compared to other simulations based on product formulas. When applied to the hydrogen chain, we have numerical evidence that our methods exhibit asymptotic scaling with the system size that is competitive with the best known qubitization approaches.

DOI: [10.1103/ynxb-p2xq](https://doi.org/10.1103/ynxb-p2xq)

I. INTRODUCTION

Among the most promising applications of quantum computing is the calculation of ground state energies in quantum many-body physics and quantum chemistry. While in general it is hard for a quantum computer to estimate the ground-state energy, quantum phase estimation (QPE) provides an efficient algorithm when a state with significant overlap to the ground state is given. The complexity of the resulting algorithms for chemistry depends largely on the complexity of simulating the Hamiltonian time evolution as a quantum circuit, as QPE requires time simulation of time $\mathcal{O}(\varepsilon^{-1})$ to compute the energy to accuracy ε . The dominant cost of performing QPE is then given by the cost of Hamiltonian evolution and its scaling with time and system size. State-of-the-art resource estimates for realistic examples from chemistry show that the computational cost of performing QPE is still formidable [1–4]. This means that for realizing useful simulations of quantum chemistry, further reductions in the circuit size required to perform Hamiltonian simulation are necessary.

This is particularly relevant for devices with a limited number of logical qubits and limited circuit depth. See Ref. [5] for a general discussion of QPE for quantum chemistry.

Deterministic product formulas, also known as Trotterization, have the advantage of being conceptually simple and to require little or no ancilla qubits for their implementation as a quantum circuit. In many situations, their performance can be better than what can be expected from rigorous bounds [1,6] or can even be proven to have good scaling, e.g., in systems with spatially local interactions [7–9]. There also exist powerful post-Trotter Hamiltonian simulation methods [10,11]; while they have better scaling, methods based on product formulas are still of great interest due to the small number of ancilla qubits required and their good performance for Hamiltonians of interest [1].

This work is motivated by quantum chemical Hamiltonians. For such a Hamiltonian, defined on an active space of N spatial orbitals (corresponding to $2N$ qubits in second quantization), the number of terms in the Hamiltonian scales as $L = \mathcal{O}(N^4)$. Trotter product formulas have a computational cost which scales with L . While this leads to polynomial scaling in N , the resulting circuits rapidly grow impractically large.

An alternative is provided by *randomized product formulas*, in particular qDRIFT [12]. Here, the number of gates no longer depends on L , but rather on λ , the sum of the interaction strengths in the Hamiltonian. If one wants to simulate Hamiltonian evolution for time t , this requires

*Contact author: f.witteveen@cwi.nl

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

$\mathcal{O}(\lambda^2 t^2)$ gates. So, while this gets rid of the dependence on L , the downside is that the scaling is quadratic in the time (and λ). In particular, since QPE with target accuracy ε requires Hamiltonian simulation for times up to $\mathcal{O}(\varepsilon^{-1})$, using randomized product formulas leads to phase estimation costs scaling with ε^{-2} rather than the optimal Heisenberg-limited scaling of ε^{-1} . Prior work uses an analysis which suggests that combining QPE with qDRIFT leads to an unfavorable scaling of the total cost with the accuracy of the computation [12]. In particular, a naive analysis leads to a scaling of order ε^{-4} or an undesirable dependence on the probability of error [4]. Using a different randomization scheme, however, as well as a different phase estimation scheme, one can restore the scaling to ε^{-2} [13].

A natural generalization of randomized product formulas is to consider *partially randomized* product formulas [14–16], which are effective if there is a relatively small number of dominant terms in the Hamiltonian, and a long tail of small terms (which are nevertheless too large to be ignored altogether). Such behavior is typical for quantum chemistry Hamiltonians; see Sec. VI and Appendix C for a detailed discussion.

A. Results and organization

We start by setting up notation for the relevant class of Hamiltonians and single-ancilla phase estimation in Sec. II and introduce background material on the Trotter formalism in Sec. III. We discuss randomized product formulas in Sec. IV. Our first contribution is an improved analysis for the cost of randomized product formulas for single-ancilla QPE, achieving reduced cost compared to previous estimates [4,12,13] and reduced maximal circuit depth. This result is stated as Theorem 1. In Fig. 1 we show the resulting resource estimates for the challenge benchmark problem FeMoco [17] and compare with previous work. We see multiple order of magnitude improvements in the performance of the randomized product formulas. Our improvement compared to Ref. [4] by more than two orders of magnitude breaks down as follows. We find a value of λ which is a factor 5 smaller, reducing the cost by a factor 25. The overhead of the phase estimation routine is reduced by a factor of 38 (while improving accuracy guarantees). Indeed, Ref. [4] finds a requirement of approximately $305\lambda^2\varepsilon^{-2}$ Pauli rotations, while we require approximately $8\lambda^2\varepsilon^{-2}$ Pauli rotations. Finally, as explained in Appendix E, we give a compilation from Pauli rotations into Toffoli gates using a factor 10 fewer gates by combining techniques from Refs. [4,13,21,22].

Figure 1 assumes an exact eigenstate, and aims to minimize total resources. More realistically, we should assume we start with a state with (high) ground state overlap. Our methods can accommodate such states, but using an exact eigenstate facilitates comparison with the previous




FIG. 1. Total Toffoli gate count and logical qubit estimates for ground-state energy estimation for FeMoco using the active space of Ref. [17], with target error $\varepsilon = 0.0016$ and assuming an initial state equal to the ground state. The blue-colored data points use qubitization [2,4,18–20], the yellow data point is the most optimistic estimate from Ref. [17] using a deterministic product formula, the red data points use randomized product formulas from Ref. [4] and this work, the orange data point uses partial randomization. We also show the Toffoli depth (i.e., the maximum number of Toffoli gates per circuit), assuming a high ground-state overlap state and choosing $\xi = 0.1$, using (partially) randomized product formulas. The vertical dashed lines connect the Toffoli depth to the total Toffoli count, to illustrate how multiple circuit runs are used.

methods. If one can prepare a state with high ground-state overlap, using robust phase estimation in combination with randomized product formulas has the effect of using multiple circuits of smaller size. This can be useful for early fault-tolerant devices and parallelization. To make this concrete, let us again consider the example of FeMoco. If one indeed can prepare exact eigenstates, or states with sufficiently high ground-state overlap, one can trade the gate count per circuit further against using more circuits. If we choose the trade-off parameter $\xi = 0.1$ in Theorem 1, we need 4741 circuits, with a maximum of 6.4×10^8 Pauli rotations per circuit (as also indicated in Fig. 1). This choice of ξ is also compatible with a state with ground-state overlap $\eta > 0.9$, and it has been argued that for a larger active space of FeMoco achieving a squared overlap of approximately 0.95 is realistic (but the cost of preparing this state is significant) [23]. As is also clear from Fig. 1, the partially randomized method (which scales better with the precision ε) benefits less from this reduced circuit depth.

The main contribution of this work is the study of partially randomized product formulas. In Sec. V we propose




FIG. 2. Gate counts for ground-state energy phase estimation for precision $\varepsilon = 0.0016$ for an N atom hydrogen chain with a minimal basis set, using deterministic, randomized, and partially randomized product formulas.

a modification of the scheme of Hagan and Wiebe [16]. We give a simple error analysis for our scheme when used in single-ancilla QPE, improving on the bounds that would be obtained by analyzing the product formula and the QPE separately. As an example of a significant cost improvement from partially randomized product formulas, we show in Fig. 2 cost estimates for the hydrogen chain with a minimal basis (a standard benchmark system to explore the scaling of the thermodynamic limit). We observe an improvement of more than an order of magnitude for partial randomization compared to fully randomized product formulas, which in turn outperform deterministic product formulas. Additionally, when varying both the required accuracy ε and the number of hydrogen atoms N , for the fully randomized method we have a scaling of $\mathcal{O}(N^{2.5}\varepsilon^{-2})$, while for the partially randomized product formula we get $\mathcal{O}(N^2\varepsilon^{-1.7})$. While the scaling with ε is suboptimal, the scaling with N is competitive with qubitization using hypertensor contraction [4], which has a scaling of $\mathcal{O}(N^{2.1}\varepsilon^{-1})$.

For electronic structure problems, we show that partial randomization can be used to improve the performance of product formulas which decompose the Hamiltonian based on a *factorization* of the Coulomb tensor [2,24]. For our quantum chemistry resource estimates, we perform a numerical estimate of the Trotter error for the hydrogen chain and a benchmark set of small molecules, and optimize of the representation of the Hamiltonian for improved performance.

Our work performs an end-to-end analysis, studying aspects of the modeling of electronic structure Hamiltonians, Hamiltonian simulation and its use as a subroutine in phase estimation. In Fig. 3 we give an overview of our main technical contributions and their relation.




FIG. 3. An overview of the topics and contributions in this work. The combined improvements lead to the resource estimates in Figs. 1 and 2.

B. Comparison to previous work

We first discuss previous work on phase estimation using randomized product formulas. Previous state-of-the-art resource estimates for QPE using qDRIFT were given in Ref. [4]. That work presents three resource analyses of phase estimation using qDRIFT, based on diamond norm bounds for the qDRIFT error. Directly requiring mean squared error ε^2 lead to a scaling of the number of gates of $\mathcal{O}(\lambda^4\varepsilon^{-4})$, which is excessively expensive. The data point on Fig. 1 is their improvement which uses a less stringent error measure, namely a 95% confidence interval of width ε . Finally, Ref. [4] gives a third approach, where the goal is to use a Hodges-Lehmann estimator on multiple samples of the QPE outcome. This reduces the phase estimation overhead by an order of magnitude, but relies on the non-rigorous assumption that the qDRIFT error is symmetric. For reference, Fig. 1 also compares with state-of-the-art estimates based on qubitization. These have better scaling in the precision and better overall performance, but do require more ancilla qubits in order to implement linear combinations of unitaries.

For single-ancilla phase estimation diamond norm bounds are not needed, and this was used in Ref. [13]. That work additionally proposed a modified randomized product formula, which we call Randomized Taylor Expansion (RTE). This has the property that it can be used to give an unbiased estimator of the time evolution unitary in single-ancilla QPE. Our analysis extends this idea to qDRIFT, which can be seen as only keeping the first order in the Taylor expansion. We also note that if one can prepare a

state with high ground-state overlap, one can break up the computation in multiple smaller circuits. We can compare our estimates for FeMoco to the ones in Ref. [13]. The algorithm in Ref. [13] is based on computing the cumulative density function of the spectrum, $C(x)$, and finding the ground-state energy as the first jump of this function. Here, we aim for error ε with failure probability at most 0.1. Assuming an exact eigenstate and finding a crossing point for $C(x) < 0.3$ or $C(x) > 0.7$, leads to an estimate of approximately 4500 circuits with an *average* number of 7.9×10^9 Pauli rotations. Here, we have used the results from Fig. 2 in Ref. [13] and rescaled to use our optimized value of λ . We see that robust phase estimation requires shorter depth circuits (without increasing the total required number).

Partial randomization schemes have also been studied previously. There are various different proposals [14–16, 25–27]. Our scheme is a modification of that proposed in Ref. [16]. The difference is that we replace qDRIFT by the RTE technique from Ref. [13] to “unbias” the product formula, which is helpful when using single-ancilla QPE. Indeed, this means that when analyzing QPE, we do not need to use diamond norm bounds (as derived in Ref. [16]), which would lead to worse constant factors. Since we find that in practice partial randomization tends to lead to (modest) constant improvements, it is important to have a sharp cost analysis.

II. SINGLE-ANCILLA QPE

We consider Hamiltonians which are a sum of local terms $H = \sum_l H_l$. A special case of interest is when

$$H = \sum_{l=1}^L H_l = \sum_{l=1}^L h_l P_l, \quad h_l \in \mathbb{R}, \quad (1)$$

where $P_l \in \pm\{I, X, Y, Z\}^{\otimes N_q}$ are Pauli operators on N_q qubits. Our main motivation derives from molecular electronic structure Hamiltonians in second quantization, as we will discuss in Sec. VI onward. In that case, we have $N_q = 2N$, where N is the number of molecular orbitals, and $L = \mathcal{O}(N^4)$. We define the *weight* of the Hamiltonian in Eq. (1) to be

$$\lambda = \sum_l |h_l|. \quad (2)$$

The eigenvalues of H and the associated eigenvectors are denoted by E_k and $|\psi_k\rangle$, respectively, where $E_0 \leq E_1 \leq E_2 \leq \dots$. Our task is to compute the ground-state energy E_0 up to precision $\varepsilon > 0$. If we know a *guiding state* $|\psi\rangle$ with significant overlap to the ground state $|\psi_0\rangle$, we can use a quantum computer to approximate E_0 by applying QPE to the time evolution operator and the initial state: $e^{-itH} |\psi\rangle$. For many problems involving the electronic




FIG. 4. The Hadamard test for estimating $Z(\theta)$. Here $R(\theta) = |0\rangle\langle 0| + e^{i\theta} |1\rangle\langle 1|$ is a phase gate.

structure of molecules, it is feasible to find good guiding states using established methods see Refs. [23, 28–32] for elaborate discussion and further references.

We consider versions of QPE based using only a single ancilla qubit, which use the *Hadamard test*, illustrated in Fig. 4. Performing QPE using only the Hadamard test (without a quantum Fourier transform circuit) has an extensive history [33–38].

We start with the guiding state $|\psi\rangle$, and a single ancilla qubit is set to $|+\rangle = H|0\rangle$. Having applied U , controlled on the ancilla qubit, we measure the ancilla with respect to the basis $|0\rangle \pm e^{i\theta} |1\rangle$. This gives a random variable Z_θ , which has outcomes ± 1 . The Hadamard test has expectation values

$$\mathbb{E}Z_\theta = \text{Re } e^{i\theta} \langle \psi | U | \psi \rangle. \quad (3)$$

In particular, letting X and Y denote the random variables obtained from $\theta = 0, \pi/2$ and setting $Z = X + iY$, we have

$$\mathbb{E}Z = \langle \psi | U | \psi \rangle. \quad (4)$$

We will focus on the case when $U = U(t) := e^{-itH}$.

If we denote by $Z_\theta(t)$ the outcomes of the Hadamard test using $U = U(t)$, then

$$\begin{aligned} g(t) &:= \mathbb{E}Z(t) = \langle \psi | \exp(-iHt) | \psi \rangle \\ &= \sum_k c_k \exp(-iE_k t), \end{aligned} \quad (5)$$

where $c_k = |\langle \psi | \psi_k \rangle|^2$. We can approximate $g(t)$ by taking multiple independent samples $Z^{(n)}(t)$ and take the mean $\bar{Z}(t) = (1/N) \sum_{n=1}^N Z^{(n)}(t)$.

If we think of $g(t)$ as a time signal, then the phase estimation routine will constitute a signal processing transformation to compute the lowest frequency of $g(t)$ (corresponding to the energy E_0) [37], provided that we have some guarantee on the overlap of $|\psi\rangle$ with the ground state; we assume a lower bound $c_0 \geq \eta$. With appropriate signal processing methods, one can find the value of E_0 with accuracy ε using m circuits with time evolution for times $t_1, \dots, t_m$. This can be done such that the maximal time evolution $t_{\max} = \max\{t_1, \dots, t_m\}$ and the total time over all circuit runs $t_{\text{tot}} = t_1 + t_2 + \dots + t_m$ both scale as ε^{-1} . This *Heisenberg scaling* is known to be optimal [39].

When using a Hamiltonian simulation method to simulate the time evolution, $t_{\max}$ is what determines the maximal number of gates per circuit and t_{tot} determines the total gate count. We measure the error of the QPE by its (root)-mean-square error. In our resource estimates we consider two different scenarios. In the first scenario, we are given access to the exact eigenstate $|\psi_0\rangle$ (so $c_0 = 1$). In the more realistic second scenario, we consider a state $|\psi\rangle$ with an overlap bound $\eta < 1$.

There are various ways to extract the ground-state energy by sampling from $\mathbf{Z}(t)$ which achieve Heisenberg scaling [36,38,40]. A particularly elegant method is *robust phase estimation* [35,41–43]. For now, we assume that the Hadamard test is implemented using *exact* time evolution, and we measure the cost in terms of the time evolution required. This method assumes that η is sufficiently large. The idea of robust phase estimation is to estimate the time evolution signal at times $t = 2^m$ for $m = 1, \dots, M$ with $M = \lceil \log \varepsilon^{-1} \rceil$. For each m , the expectation value of the Hadamard test outcome $\mathbf{Z}(2^m)$ is

$$g(2^m) = \sum_k c_k \exp(-i2^m E_k). \quad (6)$$

We repeat the test N_m times and obtain the average $\bar{\mathbf{Z}}(2^m)$. The integer M corresponds to the number of precision bits in the estimate of E_0 . With each round we update a guess θ_m for E_0 . Given an outcome $\bar{\mathbf{z}}(2^m) = r_m \exp(-i\phi_m)$ of the random variable $\bar{\mathbf{Z}}(2^m)$ and an estimate θ_{m-1} for E_0 from the previous round, the estimate θ_m is the number which is compatible with the value ϕ_m in the sense that $2^m \theta_m$ equals ϕ_m modulo 2π and is close to the estimate for θ_{m-1} from the previous round. That is, $\theta_m = 2^{-m}(\phi_m + 2\pi k)$ for the integer k such that this is closest to θ_{m-1} . It can be proven [35,41–43] that for $\eta \geq 4 - 2\sqrt{3} \approx 0.54$ this gives a correct algorithm with maximal evolution time $t_{\max} = \mathcal{O}(\varepsilon^{-1})$ and total evolution time $t_{\text{tot}} = \mathcal{O}(\varepsilon^{-1})$. In molecular electronic structure problems, in most cases it is possible to reach this ground-state overlap regime [23,28,30,31]. In situations with smaller values of η , there exist various other single-ancilla phase estimation schemes [38,40,44]. The methods in this work apply directly to these phase estimation schemes as well. The robust phase estimation scheme has the advantage that if the overlap is close to 1, one can reduce the maximal required time evolution, at the cost of increasing the total time evolution [43]. Indeed, if $\eta = 1 - \alpha$ for small α , we may take a trade-off parameter $\Omega(\alpha) = \xi \leq 1$ (i.e., we can take ξ small if the ground-state overlap is large) and obtain

$$t_{\max} = \mathcal{O}\left(\frac{\xi}{\varepsilon}\right), t_{\text{tot}} = \mathcal{O}\left(\frac{1}{\xi \varepsilon}\right). \quad (7)$$

This result is due to Ref. [43]; for completeness we provide a precise statement and proof in Theorem B1.

III. DETERMINISTIC PRODUCT FORMULAS

The previous section involved approximating e^{-iHt} as a subroutine. In this section, we explain deterministic product formulas, also known as the Trotter method, one of the leading approaches to Hamiltonian simulation, on which our work will build. We refer to Ref. [7] for an extensive discussion and references. Given a Hamiltonian $H = \sum_{l=1}^L H_l$ consisting of L terms, we may approximate for a short time evolution for time δ as

$$U(\delta) \approx S_1(\delta) := e^{-i\delta H_L} e^{-i\delta H_{L-1}} \dots e^{-i\delta H_2} e^{-i\delta H_1}. \quad (8)$$

This is a first-order product formula. A second-order product formula is

$$S_2(\delta) := \left(e^{-\frac{i}{2}\delta H_1} \dots e^{-\frac{i}{2}\delta H_L}\right) \left(e^{-\frac{i}{2}\delta H_L} \dots e^{-\frac{i}{2}\delta H_1}\right) \quad (9)$$

and higher-order Suzuki-Trotter product formulas can be recursively defined by

$$S_{2k}(\delta) = (S_{2k-2}(u_k \delta))^2 S_{2k-2}((1 - 4u_k)\delta) (S_{2k-2}(u_k \delta))^2, \quad (10)$$

for $u_k = (4 - 4^{1/(2k-1)})^{-1}$. For the Trotter-Suzuki product formula of order $2k$, $S_{2k}(\delta)$ requires $2L5^{k-1}$ unitaries of the form $e^{-i\varphi H_l}$. Compiling $e^{-i\varphi H_l}$ into elementary gates gives a quantum circuit. For example, if $H_l = h_l P_l$ are Pauli operators, $e^{-i\varphi H_l}$ is a Pauli rotation, which can straightforwardly be compiled into elementary gates.

In general, we write $S_p(\delta)$ for a p -order product formula; meaning that $S_p(\delta)$ is a product of time evolutions along the H_l . The order refers to the order of approximation in the sense that the error between $U(\delta)$ and $S_p(\delta)$ is $\mathcal{O}(\delta^{p+1})$. Time evolution for arbitrary time t can be approximated by dividing t into a sufficiently large number of intervals r and discretize $U(t)$ into r Trotter steps

$$U(t) = U(t/r)^r \approx S_p(t/r)^r, \quad (11)$$

with time step $\delta = t/r$.

We will now discuss the error due to the approximation of $U(\delta)$ by $S_p(\delta)$, which we will refer to as the *Trotter error*, in more detail. If one wants to simulate $U(\delta)$, a good error measure is the operator norm difference, which scales as

$$\|U(\delta) - S_p(\delta)\|_{\text{op}} \leq C_{\text{op}} \delta^{p+1}. \quad (12)$$

Here we have written the prefactor

$$C_{\text{op}} = C_{\text{op}}(p, \{H_l\}) \quad (13)$$

explicitly. We refer to this constant (and similar parameters, with respect to different error measures, defined

below) as the *Trotter constant*. While constant with respect to the choice of step size δ , it does depend on the order p and the Hamiltonian and its decomposition into terms H_l (and in that way on the system size). This dependence is important.

For the purpose of phase estimation for ground-state energies, we can consider a different error metric. For fixed δ , we may write $S_p(\delta) = \exp(-i\delta\tilde{H})$ for some effective Hamiltonian $\tilde{H}$ which is close to H for small δ . If we apply QPE to the unitary $S_p(\delta)$, as illustrated in Fig. 5, this is equivalent to performing QPE to the effective Hamiltonian $\tilde{H}$, and we can find an estimate for $\tilde{E}_0$, the ground-state energy of $\tilde{H}$ (note that $\tilde{H}$ depends on δ as well). The outcome of the QPE now has a *bias* due to the Trotter error, and in order to obtain an estimate of E_0 to precision ε , we need that $|E_0 - \tilde{E}_0| = \mathcal{O}(\varepsilon)$. The scaling of this bias is

$$|E_0 - \tilde{E}_0| \leq C_{\text{gs}} \delta^p \quad (14)$$

for a Trotter constant $C_{\text{gs}} = C_{\text{gs}}(p, \{H_l\})$. This means that $\delta = \mathcal{O}((C_{\text{gs}}^{-1}\varepsilon)^{1/p})$ suffices for precision ε . Phase estimation needs $\mathcal{O}(\delta^{-1}\varepsilon^{-1})$ applications of $S_p(\delta)$ to estimate $\tilde{E}_0$ to precision $\mathcal{O}(\varepsilon)$. Together, this yields a scaling of $\mathcal{O}(C_{\text{gs}}^{1/p}\varepsilon^{-1-1/p})$ Trotter steps.

The operator norm error can be bounded in the following way [7]:

$$\|U(\delta) - S_p(\delta)\|_{\text{op}} \leq \alpha_p(\{H_l\})\delta^{p+1} = \mathcal{O}((\lambda\delta)^{p+1}). \quad (15)$$

Here $\alpha_p(\{H_l\})$ is an expression in terms of operator norms of nested commutators of the terms H_l in the Hamiltonian, and λ is the weight defined in Eq. (2). One can then use these to bound

$$|E_0 - \tilde{E}_0| = \mathcal{O}(\delta^{-1} \|U(\delta) - S_p(\delta)\|_{\text{op}}) = \mathcal{O}(\lambda^{p+1}\delta^p). \quad (16)$$

The commutator bounds are often quite sharp [1,7]; a disadvantage of such bounds is that the number of commutator terms scales as $\mathcal{O}(L^{p+1})$ and may be practically infeasible to compute even for modest system sizes [45].

In Sec. VI, we numerically estimate the Trotter error constant

$$C_{\text{gs}} := \Delta E \delta^{-p-1} \quad (17)$$

from the ground-state energy error ΔE , and see that it is typically relatively small. This is an important aspect of product formula algorithms for phase estimation, since it means that Trotter methods may in practice use larger step sizes δ than one can rigorously justify and can be implemented using fewer resources than rigorous upper bounds suggest. Because of the relevance of the Trotter error for Hamiltonian simulation, this has been studied extensively

in previous work; see Refs. [6,17,45–47] for empirical and rigorous Trotter error bounds for phase estimation for quantum chemistry. Additionally, there is theoretical evidence that for most states the Trotter error is significantly lower than the worst-case bounds [48–50].

IV. PHASE ESTIMATION WITH RANDOM COMPILATION FOR TIME EVOLUTION

Two major downsides to Trotter product formulas are their scaling with L and a difficult to rigorously control Trotter error. One can avoid both these features using *randomized product formulas* [12,13,27,51–53]. Here we discuss two randomized methods of implementing an approximation to the time evolution $U(t)$: qDRIFT [12] and a randomized Taylor series expansion (RTE) [13].

We now assume that the Hamiltonian consists of a linear combination of Pauli operators P_l , since for these randomized methods we will use that $P_l^2 = I$. We will denote by

$$V_l(\phi) = e^{-i\phi P_l}, \quad (18)$$

a rotation along the Pauli operator P_l , and refer to these as Pauli rotations. Absorbing signs into the Pauli operators (i.e., replacing $P_l \mapsto \text{sgn}(h_l)P_l$), we may write the Hamiltonian as

$$H = \lambda \sum_{l=1}^L p_l P_l, \quad p_l = \frac{|h_l|}{\lambda}, \quad (19)$$

where the numbers p_l form a probability distribution p . For convenience of notation we will now assume that $\lambda = 1$; it can be reinstated by rescaling the time variable by λ . To approximate time evolution for time t , we again divide t into r steps, with a step size τ so that $\tau r = t$. The qDRIFT approximation to time evolution is given by sampling random variables $\mathbf{l}_1, \dots, \mathbf{l}_r$ according to p , and given outcomes $l_1, \dots, l_r$, apply the unitary

$$V_{l_r}(\varphi) \cdots V_{l_1}(\varphi), \quad \text{where } \varphi = \arctan(\tau). \quad (20)$$

That is, we apply a sequence of equiangular Pauli rotations, and which ones we apply is sampled proportionally to the weight of that Pauli operator in the Hamiltonian. If $\mathcal{S}_\tau$ denotes the quantum channel which with probability p_l applies $V_l(\varphi)$, the above procedure leads to the composition $\mathcal{S}_\tau^r$.

This method was proposed in Ref. [12]. In order to measure the accuracy of the approximation of the exact time evolution, we have to compare the exact time evolution unitary with the quantum channel $\mathcal{S}_\tau$. The natural measure to compare two quantum channels is the diamond norm. We let $\mathcal{U}_t$ denotes the quantum channel corresponding to applying the time evolution unitary $U(t)$ mapping

$\rho \mapsto U(t)^\dagger \rho U(t)$. It was shown in Ref. [12] that if we break the total time t into r shorter steps of time τ , so $\tau r = t$, then we have the diamond norm bound

$$\|U_t - S_\tau^r\|_\diamond \leq 2r\tau^2 \exp(2\tau) \approx 2t^2/r. \quad (21)$$

This means that choosing a number of steps $r \geq 2t^2\varepsilon^{-1}$, or reinstating the dependence on λ , $r \geq 2\lambda^2 t^2 \varepsilon^{-1}$, suffices to approximate time evolution for time t to error ε in diamond norm. In particular, there is no dependence on L , but only on λ . This is potentially beneficial if L is large and there is a large number of terms with small coefficients in the Hamiltonian.

This can then be used to simulate time evolution in QPE. Since phase estimation requires time evolution for time $t = \Theta(\varepsilon^{-1})$, this would lead to a scaling of $r = \mathcal{O}(\varepsilon^{-2})$ Pauli rotations. However, directly applying the diamond norm bounds with qDRIFT in the standard QPE algorithm leads to suboptimal results [4,12]. In general, one can ask whether using qDRIFT leads to more efficient phase estimation procedures than using Trotterization. There is no unambiguous answer: the randomized method scales worse with ε , and the relation between L and λ depends on the Hamiltonian. An advantage for qDRIFT is most likely if L is large, which is the case for molecular systems using atomic orbitals. Here, Ref. [12] found an advantage, but using very loose upper bounds for the Trotter error. In Ref. [4], for FeMoco it was found that qDRIFT was much less efficient than the Trotter estimate with optimistic Trotter error estimates from Ref. [17] (see Fig. 1).

We will now argue that previous work has overestimated the cost of doing phase estimation using qDRIFT. A key reason is that using an analysis based on diamond norm bounds is suboptimal, as also observed in Ref. [13]. Similar to the discussion of Trotter error, what we care about eventually is not an accurate simulation of the full time evolution (corresponding to small diamond norm), but rather the accuracy of the ground-state energy as extracted from the phase estimation procedure. In particular, for phase estimation methods based on the Hadamard test, what is relevant is the expectation value of the outcome of the Hadamard test. If we apply the Hadamard test to a quantum channel which applies unitary U_m with probability p_m to initial state $|\psi\rangle$, the resulting random variable $\mathbf{Z}$ has expectation value

$$\mathbb{E}\mathbf{Z} = \sum_m p_m \langle \psi | U_m | \psi \rangle. \quad (22)$$

The Hadamard test for randomized Hamiltonian simulation is illustrated in Fig. 6. The expectation value is over two sources of randomness: the sampling of the unitary U_m and the randomness from the measurement at the end of the circuit. This means that what matters for the Hadamard test is the linear combination of unitaries $\sum_k p_m U_m$ rather than

the quantum channel $\rho \mapsto \sum_m p_m U_m \rho U_m^\dagger$. Let $\mathbf{Z}(\tau, r)$ be the random variable obtained from applying a Hadamard test to r rounds of qDRIFT with interval τ . Consider the Taylor expansion of a short time evolution,

$$\begin{aligned} e^{-i\tau H} &= I - i\tau \sum_l p_l P_l + \mathcal{O}(\tau^2) \\ &= \sum_l p_l (I - i\tau P_l) + \mathcal{O}(\tau^2) \\ &= \sqrt{1 + \tau^2} \sum_l p_l V_l(\phi) + \mathcal{O}(\tau^2), \end{aligned} \quad (23)$$

where we have used that $I - i\tau P_l = \sqrt{1 + \tau^2} V_l(\phi)$ for $\phi = \arctan(\tau)$. Alternatively, we can use

$$\sum_l p_l V_l(\phi) = \frac{I - i\tau H}{\sqrt{1 + \tau^2}}. \quad (24)$$

Using this, one can show (see Appendix A for details) that, when using an initial state $|\psi\rangle$ with squared overlap c_k with the energy E_k eigenstate, the expectation value of $\mathbf{Z}(\tau, r)$ has the exact expression

$$g(\tau, r) := \mathbb{E}\mathbf{Z}(\tau, r) = \sum_k \frac{c_k}{(1 + \tau^2)^{r/2}} (I - i\tau E_k)^r. \quad (25)$$

If we choose a sufficiently small step size τ , we get

$$g(\tau, r) \approx \sum_k c_k e^{-\tau^2 r (1 - E_k^2)/2} e^{-i\tau r E_k}. \quad (26)$$

This expression is similar to $g(t)$ [the exact time signal of Eq. (5)] if we take $t = \tau r$, with the difference that there is a *damping factor* $\exp(-\tau^2 r (1 - E_k^2)/2)$. So, if we want to deduce information about the phases E_k , and we do not want to increase the number of required samples, we need to choose τ and r such that $\tau^2 r = \mathcal{O}(1)$. This leads to a scaling of $r = \Omega(t^2)$. One can reduce this value of r , but this comes at an (exponential) sampling overhead. The expectation value in Eq. (25) is an *exact* expression, so it is more amenable to analysis than the result of deterministic product formulas, where one has to bound the Trotter error. The original proposal in Ref. [12] uses an angle $\phi = \tau$; this only changes the expectation values to order $\mathcal{O}(\tau^2)$ and does not change the error analysis; here we choose $\phi = \arctan(\tau)$ for consistency with the Taylor expansion.

An alternative random product formula for Hamiltonian simulation was introduced in Ref. [13], which we call a *randomized Taylor expansion* (RTE). Whereas qDRIFT, as per Eq. (23), can be seen as a first-order expansion of $\exp(-iH\tau)$ as an linear combination of unitaries (LCU), one can also take the full Taylor expansion and normalize




FIG. 5. Performing the Hadamard test on $S_p(\delta)^s$ allows one to estimate $\langle \psi | S_p(\delta)^s | \psi \rangle = \langle \psi | \exp(-i\delta s \tilde{H}) | \psi \rangle$.

the coefficients to a probability distribution. This gives an LCU,

$$e^{-i\tau H} = B(\tau) \sum_m b_m U_m, \quad (27)$$

where the b_m form a probability distribution, and the U_m are unitaries which are a composition of one Pauli rotation and one Pauli operator and $B(\tau)$ is a normalization factor; see Appendix A for details and a derivation. Up to normalization, qDRIFT precisely corresponds to only keeping the lowest order term. The probability of sampling a higher-order term is $\mathcal{O}(\tau^2)$. The normalization factor $B(\tau)$ is approximately $\exp(\tau^2)$ for small τ and can be efficiently computed. Taking r rounds thus gives a LCU,

$$e^{-i\tau r} = B(\tau)^r \sum_{m_1, \dots, m_r} \underbrace{b_{m_1} \dots b_{m_r}}_{q_{\vec{k}}} \underbrace{U_{m_1} \dots U_{m_r}}_{W_{\vec{k}}}, \quad (28)$$

where the $q_{\vec{k}}$ are a probability distribution over $\vec{k} = (m_1, \dots, m_r)$, and the $W_{\vec{k}}$ are unitaries consisting of r Pauli rotations interspersed with Pauli operators. Performing a Hadamard test leads to a signal

$$\mathbb{E}Z(\tau, r) = \frac{1}{B(\tau)^r} \sum_k c_k e^{-iE_k \tau r}. \quad (29)$$

In particular, if we want to estimate $g(t)$, one can do so using this random variable for $\tau r = t$. The factor $B(\tau)^r \leq \exp(\tau^2 r)$ determines the required number of samples. To keep this constant, one should have $\tau^2 r = \mathcal{O}(1)$ and hence $r = \Omega(t^2)$, as for qDRIFT. The advantage of this approach is that, opposed to qDRIFT, it gives an *unbiased* estimator for the signal $g(t)$.

Since we have exact expressions for the expectation value of Hadamard tests for both qDRIFT and RTE, we can now use these to analyze single-ancilla QPE schemes. We do so for robust phase estimation. The following result is stated and proven as Theorem B1 in Appendix B and follows straightforwardly from the analysis of robust phase estimation in Refs. [35,42,43]. In contrast to the approach in Refs. [4,12], there is no dependence on a failure probability of the algorithm. We consider the case with large ground-state overlap. If $c_0 \geq \eta = 1 - \alpha$, we can choose a parameter $\xi = \Omega(\alpha)$ to reduce the maximal circuit size.

Theorem 1. Given a guiding state $|\psi\rangle$ with ground-state overlap lower bounded by $\eta \geq 4 - 2\sqrt{3} \approx 0.54$ and $\xi = \Omega(\alpha)$ for $\eta = 1 - \alpha$, robust phase estimation using random product formulas gives an estimate of E_0 with root-mean-square error ε using circuits with at most

$$r_{\max} = \mathcal{O}\left(\frac{\xi^2 \lambda^2}{\varepsilon^2}\right) \quad (30)$$

Pauli rotations per circuit and a total of

$$r_{\text{tot}} = \mathcal{O}\left(\frac{\lambda^2}{\varepsilon^2}\right) \quad (31)$$

Pauli rotations.

We also numerically estimate the constant prefactors, and for $\eta = 1$, $\xi = 1$ we find $r_{\text{tot}} \leq 8\lambda^2 \varepsilon^{-2}$ using qDRIFT. For $\eta = 1$, since we lose Heisenberg scaling, if one does not count the cost of preparing the exact ground state one might as well just measure the ground-state energy by using $\mathcal{O}(\lambda^2 \varepsilon^{-2})$ copies of the ground state. However, the robust phase estimation procedure uses only $\mathcal{O}(\log(\lambda \varepsilon^{-1})^2)$ copies of the ground state (when choosing $\xi = 1$), and crucially it still works for ground-state overlap $\eta < 1$ sufficiently large. Note additionally that the *total* cost does not depend on the trade-off parameter ξ , in contrast to the trade-off in Eq. (7). However, choosing small ξ does come at a cost of a higher number of circuits, and in particular of the number of initial state preparation circuits. For randomized methods it is advantageous to choose $\xi \sim 1 - \eta$ as small as possible (if the state preparation cost is not too high), since the total cost of the Hamiltonian simulation does not increase under the depth reduction. This is in contrast to the exact time evolution model, where a reduction in maximal evolution time comes at the cost of an increasing total evolution time, as in Eq. (7), due to the larger number of circuit repetitions. The suboptimal scaling of $\mathcal{O}(\varepsilon^{-2})$ of randomized methods is a significant disadvantage. However, as we see in molecular electronic structure examples, it can still outperform deterministic product formulas. This advantage increases further when one uses depth reductions for states with high ground-state overlap.




FIG. 6. Hadamard test for qDRIFT and for RTE. For qDRIFT, $\tau = t/r$ and $\varphi = \arctan(\tau)$. RTE is similar, but now for each gate one samples an (even) order n of the Taylor expansion. If $n = 0$, one samples l according to p as for qDRIFT. If n is larger (which happens with probability $\mathcal{O}(\tau^2)$), the rotation is over a different angle φ_n and there is an additional controlled Pauli operator in the circuit.

V. PARTIALLY RANDOMIZED PRODUCT FORMULAS

We have seen that randomization can be useful for Hamiltonian simulation in the context of phase estimation. However, it has the clear disadvantage of scaling quadratically with λ and ε^{-1} , whereas deterministic product formulas may have a better dependence on ε^{-1} . Deterministic product formulas are disadvantageous if L , the number of terms in the Hamiltonian, is large. Random product formulas are helpful if there is a tail consisting of many terms in the Hamiltonian of small weight.

A natural proposal is to treat a subset of terms in a random fashion and a subset of terms deterministically. There exist various such proposals, with different ways of implementing the partial randomization [14–16,26]. Here, we modify the scheme of Ref. [16] to a version which has the advantage of being particularly easy to analyze for the purpose of single-ancilla phase estimation. Again, the difference in the analysis is that we do not require a diamond norm bound on the resulting quantum channel, but only need to evaluate the expectation value of the resulting Hadamard test.

The partial randomization is based on a decomposition

$$H = \underbrace{\sum_{l=1}^{L_D} H_l}_{=H_D} + \underbrace{\sum_{m=1}^M h_m P_m}_{=H_R}, \quad (32)$$

where we assume that the P_m square to identity and where H_D contains the terms we treat deterministically, and H_R the terms which are treated by a random product formula. The idea is to use a deterministic Trotter formula to the decomposition of H into the terms H_l and H_R and to apply a random product formula to the decomposition of H_R into terms $h_m P_m$. The resulting circuits are shown in Fig. 7.

In order for this to yield an improvement, one aims for a decomposition such that

$$\lambda_R = \sum_m |h_m| \ll \lambda \quad \text{and} \quad L_D \ll M. \quad (33)$$

To analyze how such methods will perform, we now briefly summarize two lessons from the previous two sections. Deterministic product formulas lead to a *bias* in the ground state energy E_0 due to the Trotter error. A small step size is required in order to keep this bias below the desired precision. For randomized product formulas, one can make the resulting signal such that the ground-state energy is unbiased, but the *amplitude* of the signal is damped and one incurs a sampling overhead. Here, a sufficiently small step size is required in order to keep this sampling overhead of constant size. Partially randomized product formulas will have both a bias (from the deterministic decomposition) and a damping (from implementing the randomized part of the Hamiltonian) in the resulting signal. The error analysis straightforwardly combines these elements.

Let $S_p(\delta)$ be a deterministic product formula with respect to the decomposition of H into the terms H_l and H_R . This means that $S_p(\delta)$ is a product of unitaries of the form $e^{-i\gamma H_l}$ (which we assume to have some decomposition into elementary gates) and terms of the form $e^{-i\gamma H_R}$ for some γ . As in Eq. (14), this gives an effective Hamiltonian with ground state $\tilde{E}_0$ with error $|E_0 - \tilde{E}_0| \leq C_{\text{gs}} \delta^p$. Here, the Trotter constant $C_{\text{gs}} = C_{\text{gs}}(p, \{H_l\}_{l=1}^{L_D}, H_R)$ depends on the order p as well as on the Hamiltonian, decomposed into terms H_l for $l = 1, \dots, L_D$ and treating H_R as a single term. Each of the $e^{-i\gamma H_R}$ appearing in the formula we will then simulate using RTE, using r steps per Trotter step.

To make this concrete, we take the second-order Suzuki-Trotter formula, in which case we get $S_2(\delta)$ given by

$$e^{-\frac{1}{2}i\delta H_1} \dots e^{-\frac{1}{2}i\delta H_{L_D}} e^{-\delta i H_R} e^{-\frac{1}{2}i\delta H_{L_D}} \dots e^{-\frac{1}{2}i\delta H_1}. \quad (34)$$




FIG. 7. Hadamard test using partially randomized product formula. Here, we illustrate the second-order deterministic method, based on Eq. (34).

If we normalize H_R/λ_R , we need to do time evolution for time $\delta\lambda_R$, which we break up into r steps. Using RTE as in Eq. (28) we expand

$$e^{-i\delta H_R} = (e^{-i\tau H_R/\lambda_R})^r \quad \text{for } \tau r = \delta\lambda_R \quad (35)$$

$$= B \sum_k q_k W_k \quad (36)$$

as a LCU, which gives a LCU for $S_2(\delta)$. Here, each W_k is a composition of r Pauli rotations and Pauli operators.

Since this gives an unbiased LCU for the exact time evolution, it is straightforward to show (see Lemma A3) that using RTE and repeating s times in a Hadamard test has expectation value

$$\mathbb{E}Z(\delta, r, s) = \frac{1}{B(r, \delta)^s} \langle \psi | S_p(\delta)^s | \psi \rangle, \quad (37)$$

where

$$B(r, \delta) \leq \exp\left(\mathcal{O}\left(\frac{\lambda_R^2 \delta^2 s}{r}\right)\right). \quad (38)$$

In order to perform phase estimation, we need to choose δ such that the Trotter error of $S_p(\delta)$ is smaller than ε . That means we should choose a Trotter step size $\delta = \mathcal{O}((C_{\text{gs}}\varepsilon)^{\frac{1}{p}})$ and maximal depth $s = \mathcal{O}(C_{\text{gs}}^{\frac{1}{p}}\varepsilon^{-1-\frac{1}{p}})$ for phase estimation. In order to keep the number of samples constant that are required to estimate $\langle \psi | S_p(\delta)^s | \psi \rangle$ to constant precision, we need

$$B(r, \delta)^s \leq \exp\left(\mathcal{O}\left(\frac{\lambda_R^2 \delta^2 s}{r}\right)\right) \quad (39)$$

to be constant. Since $\delta s = \mathcal{O}(\varepsilon^{-1})$, this can be achieved by taking the total number of rotations in the randomized part

as $rs = \mathcal{O}(\lambda_R^2 \varepsilon^{-2})$. Together, this leads to the following result.

Theorem 2. Let H as in Eq. (33), and suppose that the Trotter error constant for H (written as a sum of the H_l and treating H_R as a single term) is C_{gs} , then phase estimation to precision ε using the partially randomized product formula requires $\mathcal{O}(L_D C_{\text{gs}}^{\frac{1}{p}} \varepsilon^{-1-\frac{1}{p}})$ evolutions along the H_l and $\mathcal{O}(\lambda_R^2 \varepsilon^{-2})$ Pauli rotations.

Gate counts with constant factors are given in Appendix E. A subtlety here is that the Trotter constant C_{gs} depends on the choice of partitioning. This complicates the task of finding an optimal distribution of terms to H_R or H_D . One can bound C_{gs} using a commutator bound in the usual way (see Ref. [16] for such a bound). For our numerical estimates we will take a different approach: we estimate the Trotter error for the fully deterministic method and use the resulting Trotter error constant. In that case, it is clear that to minimize the cost, the optimal decomposition assigns the L_D largest weight terms to H_D for some choice of L_D , and the remainder to H_R . Suppose that H has L terms and weight λ . Clearly, partial randomization is most useful if there is a small number $L_D \ll L$ of terms with large weight (so including these in H_D means that $\lambda_R \ll \lambda$); see Fig. 8 for an example of a distribution of weights for a quantum chemistry Hamiltonian.

VI. ELECTRONIC STRUCTURE HAMILTONIANS

We start with a molecular electronic structure Hamiltonian in second quantization,

$$\sum_{pq\sigma} h_{pq} a_{p\sigma}^\dagger a_{q\sigma} + \frac{1}{2} \sum_{pqrs, \sigma\tau} h_{pqrs} a_{p\sigma}^\dagger a_{r\tau}^\dagger a_{s\tau} a_{q\sigma}, \quad (40)$$




FIG. 8. Distribution of the weights of terms of the FeMoco Hamiltonian in the Pauli decomposition (top) and the 1- and 2-electron terms in the double-factorization representation (bottom). The dashed lines show an exponential fit to the tail (main plot) and a power law for the largest terms (the log-log inset plot).

where $a_{p\sigma}^\dagger$, $a_{q\sigma}$ are fermionic creation and annihilation operators and h_{pq} , h_{pqrs} are integrals arising from a calculation using the underlying orbital basis (see Appendix C). The indices p, q, r, s denote the spatial degrees of freedom, whereas σ, τ are reserved for spin. The Hamiltonian acts on the space of $2N$ spin orbitals, and we fix the number of electrons n . And a choice of fermion-to-qubit mapping transforms it into a linear combination of Pauli operators, acting on $2N$ qubits with $L = \mathcal{O}(N^4)$ terms gives a representation of H as a linear combination of Pauli operators

$$H = \sum_l H_l = \sum_l h_l P_l. \quad (41)$$

We work with active spaces of molecular orbitals. An alternative choice is a plane wave basis, which has the advantage that the Hamiltonian has fewer terms and the Trotter error can be more easily controlled. This is particularly useful for materials [7–9, 47, 54]. For this work we

restrict to active spaces based on molecular orbitals since plane wave bases need a much larger number of orbitals to reach similar accuracy. While in principle there are $\mathcal{O}(N^4)$ nonzero coefficients in Eq. (40), many are extremely small and can be truncated. For large spatially extended systems, one expects that the number of significant terms mostly scales quadratically in N . Generally, one expects polynomial decay of the weight, with a tail of terms which decays exponentially as in Fig. 8. See Appendix C for more discussion.

As benchmark systems, we use a linear chain of hydrogen atoms, which is a toy model for a system of increasing size, and FeMoco, a standard challenge benchmark problem used for quantum computing resource estimates, as well as a collection of small molecules with varying active space sizes. Detailed descriptions can be found in Appendix C.

The computational cost of implementing product formulas for H depends on various factors: for deterministic product formulas, it depends on L and the Trotter error; for randomized product formulas it depends on the weight λ , and for partially randomized methods it additionally depends on a choice of partitioning into deterministic and randomized terms. In order to assess the performance of phase estimation with partially randomized product formulas, in the following subsections we discuss numerical estimates of the Trotter error, and the decomposition into H_D and H_R .

A. Estimating Trotter error

In order to give a good comparison between deterministic and (partially) randomized product formulas, we start by evaluating the Trotter error, which is crucial for the performance of the deterministic product formula. We perform a detailed analysis of Trotter error with system size.

Apart from the practical difficulties in evaluating the error bounds based on commutators in Eq. (15), for the purpose of computing ground-state energies bounding the ground-state energy error in terms of the operator norm error can be a loose bound. This has been confirmed by numerical chemistry simulations for small molecules [6]. From a theoretical perspective, the operator norm of our matrices depends on how many orbitals we choose to include, which is a choice of the algorithm and not a fundamental physical feature of the problem; see Ref. [55] for more discussion of this.

We numerically estimate the Trotter error constant for ground-state energy errors, and see that it is typically relative small, as shown in Fig. 9. See also Refs. [6, 17, 45–47] for previous work on empirical and rigorous Trotter error bounds for phase estimation for quantum chemistry. This is an important aspect of product formula algorithms for phase estimation, since it means that Trotter methods may




FIG. 9. We calculate the Trotter error ΔE (on the ground-state energy) and use this to estimate the prefactor C_{gs} , following Eq. (17). This is done for $p = 2$ for a variety of molecules and active spaces, and plotted against λ .

in practice use larger step sizes δ than we can rigorously justify and can be implemented using fewer resources than rigorous upper bounds suggest.

As discussed in Sec. III, there are various error measures and bounds one can use. For our purposes, we care about the ground-state energy error. In order to estimate this error, we consider a benchmark set of small molecules with varying active space size and at half filling, so the number of electrons n equals half the number of spin-orbitals $2N$. We compute the Trotter error constant numerically for the second order method with $p = 2$. As we show in Fig. 9, the Trotter error constant correlates well with λ . Crucially, even though the Trotter error constant scales with system size, it is typically rather small even for larger systems. These results corroborate similar findings in Refs. [6,17].

Further numerical results are presented in Appendix D. We estimate Trotter error scalings for the hydrogen chain, which has a natural scaling limit. In general we see that using the ground-state energy error instead of an operator norm bound on the full unitary leads to about an order of magnitude improvement in the cost estimate. Additionally, for our chemistry benchmark systems we see no structural improvement by going to higher-order Trotter-Suzuki product formulas.

B. Choice of decomposition for partial randomization in the Pauli representation

For the partially randomized product formulas, one additionally needs to choose a partition of the Hamiltonian into a deterministic and randomized part as in Eq. (32). We first discuss the case where we consider H in a Pauli representation, where the terms H_D and H_R are all Pauli operators [i.e., $H_I = P_I$ in Eq. (32)]. The total cost of the

partially randomized product formula circuit is the sum of the deterministic and randomized parts as in Theorem 2. Obviously, we want to choose a decomposition which minimizes this sum.

A subtlety is that the Trotter error of the scheme depends on the choice of decomposition, and which terms are treated randomly; i.e., $C_{\text{gs}} = C_{\text{gs}}(p, \{H_I\}, H_R)$. A (rigorous) approach to this is to use commutator bounds to estimate the Trotter error and to take this into account in the optimization problem. Such an approach is pursued in Ref. [16] and applies straightforwardly to our method as well. However, this leads to a complicated optimization problem. Additionally, for the type of problems we consider, the deterministic methods need to be taken with heuristic empirical error scaling in order to be competitive. As a heuristic, we simply bound the Trotter error of the partially randomized method by our estimate of the Trotter error of the fully deterministic method. This is not a rigorous treatment but appears to be reasonable in practice. We study the dependence of the Trotter error on the choice of decomposition numerically. The results can be found in Appendix D; the conclusion is that when decreasing L_D , the Trotter error does not decrease significantly unless λ_R is a significant fraction of λ . This means that in practice the heuristic of using the Trotter error of the fully deterministic method should suffice; it is not valid when λ_R is close to λ but in this regime partially randomized product formulas will anyways not improve much over fully randomized product formulas.

Once one assumes that the Trotter error constant C_{gs} may be taken constant when finding the optimal decomposition, it is clear that the optimal decomposition should simply include the L_D highest weight terms in H_D , and the remainder in H_R , for some choice of L_D . One can then easily numerically search for the optimal value of L_D which minimizes the cost in Theorem 2. In Fig. 2 we compare the gate counts for phase estimation of the hydrogen chain, comparing the deterministic, randomized (qDRIFT) and partially randomized methods. In terms of cost metric, we consider the number of non-Clifford gates (expressed as number of Toffoli gates) after compilation of the Pauli rotations. We also count the number of 2-qubit gates, see Appendix E. In both cases, the partially randomized method achieves a reduction of cost, while scaling similar to the randomized method with increasing system size. We note that the optimal choice of decomposition depends on the cost metric employed (e.g., whether we measure two-qubit gates or non-Clifford gates).

Another interesting use for partial randomization is that one can improve the compilation to elementary gates. The randomized part consists of many rotations over the same angle, which is beneficial when compiling to Toffoli gates [47,56]. For the deterministic part $H_D = \sum_I h_I P_I$, one can round the values of the h_I to values more favorable for compilation by transferring a small part of the weight to the

randomized part, see Appendix E. This is similar to previously known randomization strategies for compilation [57–59].

Finally, one could ask to what extent the terms in H_R contribute to the value of the ground-state energy. In principle, it could be that one can deterministically truncate many terms in the Hamiltonian without perturbing the ground-state energy by much. We study this numerically in Appendix D, and find that at the partition with optimal cost for partial randomization, the error of fully truncating H_R is generally too large.

VII. DIFFERENT REPRESENTATIONS OF THE HAMILTONIAN

The representation in Eq. (41) is not unique. A significant degree of freedom is the choice of orbital basis. In Sec. VII A we discuss how this degree of freedom can be optimized to improve the performance of randomized product formulas. Moreover, for deterministic product formulas it can be beneficial to use different decompositions of the Hamiltonian based on a factorization of the Coulomb tensor. In Sec. VII B we argue that this approach benefits from partial randomization as well.

A. Weight reduction for the Hamiltonian

The performance of randomized product formulas depends on the value of the weight λ . Therefore, minimizing the value of λ is important for (partially) randomized product formulas. It is in fact relevant for all methods based on LCU, such as qubitization. For this reason, it has been investigated extensively [60,61]. We employ two techniques to significantly reduce the value of λ .

The first is that one can choose a different basis for the single-particle Hilbert space (i.e., a unitary transformation of the orbitals) [60]. For choosing a good basis of orbitals, a good starting point are localized orbitals, rather than the canonical orbitals derived from a Hartree-Fock calculation. While this choice already significantly reduces the value of λ , one can also use these as a starting point for further minimization [60]. One can directly optimize λ over parameterized orbital basis transformations via gradient descent. Since this is a high-dimensional optimization problem it is crucial to start from a good initial guess. Typically, using localized orbitals rather than the canonical orbitals derived from a Hartree-Fock calculation reduces the value of λ . In case orbital localization does not converge or does not yield a lower λ value, we find that the orbital rotation into the eigenbasis of the largest Cholesky vector of the Coulomb-tensor decomposition to be an inexpensive alternative. Besides the choice of orbitals, the Hamiltonian has a particle-number symmetry which can be exploited, and this leads to a second technique to find a different Hamiltonian with the same ground-state energy, but with lower weight [62].

Applying these methods to FeMoco for the Reiher [17] and Li [63] choice of active space, we see a reduction of the value of λ from 1863 to 405 for the Reiher Hamiltonian and a reduction from 1511 to 568 for the Li Hamiltonian. We note that this also leads to a significant improvement of the performance of the sparse qubitization method of Ref. [18] to FeMoco. This leads to the results reported in Fig. 1; see Appendix E for details.

When we discussed the choice of decomposition for the partially randomized method, we assumed a fixed choice of orbitals and hence of the coefficients in the Pauli representation of the Hamiltonian. A reasonable choice is to take an orbital basis for which λ is small. However, a more complicated optimization is possible: one would want to find an orbital basis where most of the weight is concentrated in a small number of terms. We leave such optimization to future work.

B. Partial randomization with factorizations of the Hamiltonian

For deterministic product formulas, one does not need to restrict to the Pauli representation of the Hamiltonian. One general approach is based on factorizations of the Hamiltonian, where one can factorize the Coulomb tensor h_{pqrs} . This gives a representation of the two-body terms of the Hamiltonian as

$$\sum_{l=1}^L H_l = \sum_{l=1}^L (U^{(l)})^\dagger \left(\underbrace{\sum_{p=1,\sigma}^{\rho_l} \lambda_p^{(l)} n_{p\sigma}}_{V^{(l)}} \right)^2 U^{(l)}, \quad (42)$$

where for each l we have a different orbital basis, $U^{(l)}$ is the unitary implementing the change of orbital basis. See Appendix C for details. The $n_{p\sigma} = a_{p\sigma}^\dagger a_{p\sigma}$ are the number operators. After a Jordan-Wigner transformation, the Hamiltonian $V^{(l)}$ is diagonal in the standard basis. This can be used for a deterministic product formula [24], since one can implement $\exp(-i\delta H_l)$ by

$$\exp(-i\delta H_l) = (U^{(l)})^\dagger \exp(-iV^{(l)}) U^{(l)}. \quad (43)$$

Both the orbital basis changes and the diagonal Hamiltonian evolutions can be performed using $\mathcal{O}(N^2)$ gates using Givens rotations [64], so in total a single Trotter step needs $\mathcal{O}(LN^2)$ gates (both in terms of number of two-qubit gates and single-qubit rotations).

A decomposition of the form in Eq. (42) can be obtained through a *double factorization*, which we review in Appendix C. Here, the first factorization is into a sum of L terms, the second factorization is the diagonalization of these in different orbital bases. This can be an improvement over the Pauli representation if the ranks L and ρ_l are small. Decompositions of the Hamiltonian as in Eq. (42),

or similar decompositions are also essential to qubitization and other LCU based approaches to Hamiltonian simulation [2,4,18]. Typically, in factorized representations, one truncates a tail of terms which are so small that they do significantly affect the ground-state energy. Partial randomization allows for much more flexible truncation schemes, and in that way can be used to improve product formulas based on Hamiltonian factorizations.

First of all, in Eq. (42), some of the terms H_l are rather small (but not so small that they can be completely ignored). As an example, the distribution of the sizes of the H_l for a double factorization of the FeMoco Hamiltonian is shown in Fig. 8. This shows, similar to the Pauli representation, a dominant set of terms obeying powerlaw decay, and a tail of terms with exponential decay. See Ref. [65] for a discussion of the scaling of ρ_l for spatially extended systems. Note that the number of terms is much smaller than in the Pauli representation [but the cost of implementing any of the $\exp(-i\delta H_l)$ is larger]. We can apply partial randomization by keeping the L_D largest terms in the deterministic part. The remainder term is now not given as a linear combination of Paulis, but we can easily re-express the remainder terms $\sum_{l>L_D} H_l$ in terms of Pauli operators using some choice of orbital basis. In other words, we may use the partial randomization to truncate the rank L of the first factorization.

One of the main bottlenecks of the factorization approach to Trotterization is the cost of implementing basis changes on the quantum computer, which scales quadratic in the number of orbitals. This cost can be reduced if the diagonal Hamiltonian $\sum_{p\sigma} \lambda_p^{(l)} n_{p\sigma}$ has only $\rho_l < N$ nonzero terms. In that case, the basis transformation can be arbitrary on the remaining $N - \rho_l$ orbitals. Such basis change unitaries can be implemented using $\mathcal{O}(N\rho_l)$ gates [64]. It turns out that for large systems, for each l , the λ_p can be truncated to ρ_l terms, where ρ_l is significantly smaller than N . Truncating small λ_p leads to constant ρ_l for large N ; this effect, however, kicks in at rather large N [65,66]. We can again use partial randomization to truncate the second factorization. Indeed, in this case for the diagonal Hamiltonian $V^{(l)}$ we only keep the terms for a subset of ρ_l orbitals for which $\lambda_p^{(l)}$ is relatively large. The remainder is assigned to the randomized part H_R . In Fig. 10 we show that this leads to a significant improvement for the hydrogen chain. Additionally, it resolves the issue that for the deterministic approach it is ambiguous what are good truncation thresholds which do not perturb the ground-state energy too much. We note that for the hydrogen chain, the overall gate count using the Pauli representation, as shown in Fig. 2 is lower.

As a final comment, we note that for the randomized methods we assume the terms in the Hamiltonian are Pauli operators. However, for the derivations the only relevant property is that $P_l^2 = I$. This means that one can also use a Hamiltonian of the form $H = \sum_m h_m \tilde{P}_m$ where




FIG. 10. Toffoli gate counts for phase estimation of the hydrogen chain with N atoms using a factorized representation of the Hamiltonian.

each $\tilde{P}_m$ is a Pauli with respect to a different choice of orbital basis. For example, one obtains such representations from the double factorization representation of the Hamiltonian. In that case, $V^{(l)}$ is a sum of terms proportional to $n_{p\sigma} n_{q\tau}$, which under the Jordan-Wigner mapping can be expanded as Pauli operators acting on only two qubits. Therefore, the Hamiltonian can be expressed as a sum of terms of the form $\tilde{P}_m = (U^{(l)})^\dagger P_m U^{(l)}$, where $U^{(l)}$ is an orbital basis change unitary and P_m is a two-qubit Pauli operator. The cost of implementing $\exp(-i\varphi \tilde{P}_m) = (U^{(l)})^\dagger \exp(-i\varphi P_m) U^{(l)}$ is dominated by the cost of implementing the basis change, and since P_m only acts on two qubits, this cost is $\mathcal{O}(N)$. So, in cases where such a representation drastically reduces the weight λ of H , this may be a beneficial approach for randomized product formulas. Let λ be the weight of H in a Pauli representation, and let λ_{DF} be the weight in a double factorization representation. Since the cost scales quadratically in the value of λ , we have an improvement if $\lambda/\lambda_{\text{DF}} = o(\sqrt{N})$. As an example, Ref. [4] reports for an arrangement for 4 hydrogen atoms with an increasing number of orbitals (as a toy model for the continuum limit), scalings of $\lambda = \mathcal{O}(N^{3.1})$ for the Pauli representation and $\lambda_{\text{DF}} = \mathcal{O}(N^{2.3})$, so this approach would asymptotically improve the performance of random product formulas.

VIII. CONCLUSION

We performed a comprehensive analysis of QPE using the Hadamard test and time evolution by product formulas for deterministic and (partially) randomized product formulas. Deterministic product formulas are conceptually simple and their scaling with the approximation error ε is close to optimal when using higher-order methods. The main downside is their scaling with the number of terms in the Hamiltonian L , and a potentially complicated dependence on the Trotter error. Random product formulas have

no dependence on the number terms, but rather on the weight λ . However, the scaling with ε^{-1} and λ is quadratic, which is suboptimal. In this work, we analyze phase estimation and Hamiltonian simulation in an integrated way. In particular, this allows us to obtain a sharper analysis of phase estimation for random product formulas. We also prove that the maximal depth can be reduced from ε^{-2} to $(1 - \eta)^2 \varepsilon^{-2}$ in regimes with large ground-state overlap. These improvements do not increase the total number of gates, in contrast to other methods (which have better scaling with ε).

Our other main contribution is a modification of the partially randomized scheme of Ref. [16] and an analysis of its use in single-ancilla phase estimation. This approach combines advantages of deterministic and randomized product formulas. We demonstrated that this can lead to an improvement in the required resources for ground-state energy estimation.

For a wide variety for benchmark molecular systems we compare the cost of different product formula Hamiltonian simulation methods for phase estimation. For deterministic product formulas, it is well known that (computable) rigorous upper bounds can heavily overestimate the Trotter error, leading to much shorter Trotter steps (and hence deeper circuits) than necessary. We numerically analyze the empirical Trotter error, and show that while this gives a large improvement in the required resources, random product formulas can still outperform deterministic product formulas. Partial randomization can lead to a further reduction in cost.

On a high level the decomposition strategies into deterministic and randomized parts H_D and H_R are guided by the following aim: Find an approximation H_D to the full Hamiltonian H , such that the time evolution under H_D is easier to implement than for H , while the remainder is small, $\lambda_R < \lambda$ (which therefore can be implemented cheaply in a randomized way). This is a change of perspective compared to previous works on optimized Hamiltonian representations. The possible approximations H_D of the electronic structure Hamiltonian allow for a large degree of freedom in the algorithm design. We proposed a decomposition strategy based on factorizations of the Coulomb tensor, which lead to improved resource estimates compared to the fully deterministic method. One can further optimize the representation of H by choosing different orbital bases for the deterministic and randomized Hamiltonians or use the structure of the Coulomb tensor to optimize the cost. We think this is a promising direction for future explorations.

In this work we are counting both non-Clifford gates as well as two-qubit gates. It is unclear whether in practice the non-Clifford gates are necessarily the dominant cost in fault-tolerant quantum computation [56,67]. In particular, it appears that magic state distillation may be cheaper for logical error rates around 10^{-9} [56], which means

that trade-offs which reduce circuit depth may be extra useful. Product formulas (using many very small angle single-qubit rotations) are also a promising candidate for approaches which are only partially fault-tolerant [68]. The quantum circuits deriving from product formulas have the advantage of being structurally simple (they consist only of Pauli rotations if the Hamiltonian is a linear combination of Paulis), so it is straightforward to count or optimize other cost measures.

Together with various other developments in the algorithmic theory of product formulas, such as multiproduct formulas and polynomial interpolation techniques [69–72], isometries to enlarged basis sets [73], this suggests that for both intermediate-term applications and fully fault-tolerant quantum computation (partially randomized), product formulas remain a promising approach to ground-state energy phase estimation.

ACKNOWLEDGMENTS

We thank Li Liu for helpful discussions. We are grateful to Markus Reiher for providing access to Euler supercomputer at ETH Zürich. We acknowledge funding from the Research Project “Molecular Recognition from Quantum Computing.” Work on “Molecular Recognition from Quantum Computing” was supported by Wellcome Leap as part of the Quantum for Bio (Q4Bio) Program. We also acknowledge financial support from the Novo Nordisk Foundation (Grant No. NNF20OC0059939 “Quantum for Life”).

DATA AVAILABILITY

The data that support the findings of this article are openly available [74].

APPENDIX A: HADAMARD TEST USING (PARTIALLY) RANDOMIZED TIME EVOLUTION

In this appendix we derive in detail the result of performing a Hadamard test on (partially) randomized product formulas. We briefly recall the setup and notation. The Hamiltonian is given as

$$H = \sum_l H_l, \quad (\text{A1})$$

where for the randomized product formulas we will assume more specifically that

$$H = \sum_l h_l P_l, \quad (\text{A2})$$

where the P_l are Pauli operators and $h_l \geq 0$ (using that signs can be absorbed in the Pauli operators). In fact, the only properties used in the derivations are the possibility to

implement time evolution along P_{l_i} and $P_l^2 = I$. We write $V(\varphi) = \exp(-i\varphi P_l)$. We may rescale the Hamiltonian by a factor

$$\lambda = \sum_l h_l \quad \text{so} \quad H = \lambda \sum_l p_l H_l, \quad (\text{A3})$$

where the p_l form a probability distribution p . We may rescale time by a factor λ , and thereby assume that $H = \sum_l p_l P_l$ where the p_l form a probability distribution p . In the end, when determining the computational cost we have to account for the rescaling by rescaling the relevant parameters by λ . The Hamiltonian has energies E_k (which lie in $[-1, 1]$ after rescaling with λ) with eigenprojectors Π_k . We apply the Hadamard test with initial state $|\psi\rangle$, which has squared overlap $c_k = \langle \psi | \Pi_k | \psi \rangle$ with the eigenspace of energy E_k . If E_k is nondegenerate, $\Pi_k = |\psi_k\rangle\langle\psi_k|$ and $c_k = |\langle \psi | \psi_k \rangle|^2$.

1. Hadamard test using qDRIFT

The qDRIFT method is to construct a unitary

$$U(\tau, r) = \prod_{i=1}^r \exp(-i\varphi H_{l_i}), \quad (\text{A4})$$

where l_i is sampled according to p and φ depends on a choice of time step τ . This gives a random unitary, but it turns out that for appropriately chosen r and φ the channel obtained by this process is close in diamond norm to the time evolution unitary $U(t) = \exp(-iHt)$ [12,51]. This may then be used in combination with phase estimation to compute ground-state energies.

In this work we analyze in detail the performance of phase estimation methods based on the Hadamard test using qDRIFT. The analysis is very similar to that of Ref. [13], which also uses Hadamard tests together with a randomized method for time evolution. The main difference is that in their approach the randomized time evolution method is based on a Taylor series expansion of the time evolution operator.

The fact that P_l squares to identity implies that for any $\varphi \in \mathbb{R}$ the operator

$$V_l(\varphi) := \cos(\varphi)I - i\sin(\varphi)P_l = \exp(-i\varphi P_l) \quad (\text{A5})$$

is unitary. We now choose $\tau \in [0, 1]$ and let $\varphi = \arctan(\tau)$. Let $X(\tau, r)$ and $Y(\tau, r)$ be the ± 1 -valued random variables resulting from the following process:

- (1) Sample $l_1, \dots, l_r$ from p .
- (2) Apply a Hadamard test to the unitary

$$U = V_{l_r}(\varphi) \cdots V_{l_1}(\varphi). \quad (\text{A6})$$

Let $Z(\tau, r) = X(\tau, r) + iY(\tau, r)$. We may now compute the expectation value of $Z(\tau, r)$, which we denote by $g(\tau, r)$ to be the following signal:

$$g(\tau, r) := \mathbb{E}Z(\tau, r) \quad (\text{A7})$$

$$= \sum_{l_1, \dots, l_r} p_{l_1} \cdots p_{l_r} \langle \psi | V_{l_r}(\varphi) \cdots V_{l_1}(\varphi) | \psi \rangle \quad (\text{A8})$$

$$= \sum_{l_1, \dots, l_r} p_{l_1} \cdots p_{l_r} \langle \psi | \frac{I - i\tau P_{l_r}}{\sqrt{1 + \tau^2}} \cdots \times \frac{I - i\tau P_{l_1}}{\sqrt{1 + \tau^2}} | \psi \rangle \quad (\text{A9})$$

$$= (1 + \tau^2)^{-r/2} \langle \psi | (I - i\tau \sum_{l_r} p_{l_r} P_{l_r}) \cdots (I - i\tau \sum_{l_1} p_{l_1} P_{l_1}) | \psi \rangle \quad (\text{A10})$$

$$= (1 + \tau^2)^{-r/2} \langle \psi | (I - i\tau \sum_l p_l P_l)^r | \psi \rangle \quad (\text{A11})$$

$$= (1 + \tau^2)^{-r/2} \langle \psi | (I - i\tau H)^r | \psi \rangle, \quad (\text{A12})$$

which encodes the spectrum of the (normalized) Hamiltonian. We summarize this conclusion in the following lemma.

Lemma A1. Let

$$H = \sum_l p_l P_l, \quad (\text{A13})$$

where $P_l^2 = I$ and the p_l are a probability distribution. Then the random variables Z from the measurement outcomes of a Hadamard test using qDRIFT with r Pauli rotations and step size τ has expectation value,

$$\mathbb{E}Z(\tau, r) = g(\tau, r) = (1 + \tau^2)^{-r/2} \langle \psi | (I - i\tau H)^r | \psi \rangle. \quad (\text{A14})$$

This can be expressed in terms of the eigenstate overlaps and energies as

$$\begin{aligned} g(\tau, r) &= (1 + \tau^2)^{-r/2} \sum_k c_k (1 - i\tau E_k)^r \\ &= \sum_k c_k \underbrace{\left(\frac{1 + E_k^2 \tau^2}{1 + \tau^2} \right)^{\frac{r}{2}}}_{\geq \exp(-\frac{1}{2}(1 - E_k^2)r\tau^2)} \exp[-ir \arctan(\tau E_k)]. \end{aligned} \quad (\text{A15})$$

The goal is to extract the value of E_0 from this signal, choosing appropriate values of τ and r . Recall the signal obtained in Eq. (5) where we perform time evolution for

time t . If we choose $r = \lceil t^2 \rceil$ and $\tau = t/r \approx t^{-1}$ we find that

$$g(\tau, r) = \sum_k d_k e^{-itF_k}, \quad (\text{A16})$$

with

$$F_k = \tau^{-1} \arctan(\tau E_k) \approx tE_k, \quad c_k \exp\left(-\frac{1}{2}(1 - E_k^2)\right) \leq d_k \leq c_k, \quad (\text{A17})$$

which approximates $g(t)$, with the difference that the amplitudes are modified by a factor at most $\sqrt{e}$.

2. Randomized Taylor expansion

In this appendix we review a refinement of qDRIFT due to Berta, Campbell and Wan [13], which we will call the *Randomized Taylor Expansion* (RTE) method. It proceeds in very similar fashion to qDRIFT, and in fact the above qDRIFT algorithm corresponds exactly to keeping only the first-order term in the Taylor expansion. We write a Taylor series expansion for $\exp(-i\tau H)$ from which we get

$$\exp(-i\tau H) = \sum_{n=0}^{\infty} \frac{(-i\tau)^n}{n!} \left(\sum_l p_l P_l \right)^n \quad (\text{A18})$$

$$= \sum_{n \text{ even}} \frac{(-1)^{n/2}}{n!} \tau^n \left(I - \frac{i\tau}{n} \sum_l p_l P_l \right) \left(\sum_l p_l P_l \right)^n \quad (\text{A19})$$

$$= \sum_{n \text{ even}} \frac{(-1)^{n/2}}{n!} \tau^n \sum_l p_l \left(I - \frac{i\tau}{n+1} P_l \right) \left(\sum_l p_l P_l \right)^n \quad (\text{A20})$$

$$= \sum_{n \text{ even}} \frac{(-1)^{n/2}}{n!} \tau^n \sqrt{1 + \frac{\tau^2}{(n+1)^2}} \sum_l p_l V_l(\varphi_n) \times \left(\sum_l p_l P_l \right)^n, \quad (\text{A21})$$

where $\varphi_n = -\arctan(\tau/(n+1))$. This gives a decomposition

$$\exp(-i\tau H) = \sum_m c_m U_m, \quad (\text{A22})$$

where m is a label consisting of an even integer n and $l, l_1, \dots, l_n$,

$$c_m = \frac{1}{n!} \tau^n \sqrt{1 + \frac{\tau^2}{(n+1)^2}} p_l p_{l_1} \dots p_{l_n} \geq 0, \quad (\text{A23})$$

$$U_m = (-1)^{n/2} V_l(\varphi_n) P_{l_n} \dots P_{l_1}.$$

Note that in principle, n can be arbitrary, but one can truncate to finite n at small error, see Ref. [13]. Again, for simulation of time t , we can break up in r steps, and for $\tau = t/r$,

$$\exp(-i\tau H)^r = \sum_{m_1, \dots, m_r} c_{m_1} \dots c_{m_r} U_{m_r} \dots U_{m_1} \quad (\text{A24})$$

$$= \sum_k b_k W_k, \quad (\text{A25})$$

where k is a label consisting of $m_1, \dots, m_r$ and $b_k = c_{m_1} \dots c_{m_r}$, and the unitary W_k is given by $U_{m_r} \dots U_{m_1}$. If we now let

$$B = \sum_k b_k \quad q_k = \frac{b_k}{B}, \quad (\text{A26})$$

then sampling according to q_k and applying a Hadamard test to W_k with initial state $|\psi\rangle$ gives random variables $X(\tau, r)$ and $Y(\tau, r)$ such that for $Z(\tau, r) = X(\tau, r) + iY(\tau, r)$,

$$\mathbb{E}Z(\tau, r) = \frac{1}{B} \langle \psi | \exp(-itH) | \psi \rangle. \quad (\text{A27})$$

This means that we can estimate the value of

$$g(\tau r) = g(t) = \langle \psi | \exp(-itH) | \psi \rangle \quad (\text{A28})$$

to precision γ by estimating the expectation value of $Z(\tau, r)$ to precision γ/B . In other words, if we want to compute $g(\tau, r)$ using this method we can do so at a sampling overhead of B^2/γ^2 . One can bound [12]

$$B = \sum_k b_k = \left(\sum_m c_m \right)^r \leq \exp(\tau^2 r). \quad (\text{A29})$$

This means that one can estimate $g(t)$ with constant sampling overhead by taking the same choice of parameters as for qDRIFT with depth $r = \lceil t^2 \rceil$ and step size $\tau = t/r \approx t^{-1}$.

Given a sample $k = (m_1, \dots, m_r)$ the resulting unitary W_k consists of the composition of r unitaries U_{m_i} , which itself consists of a single Pauli rotation together with a product of Pauli operators (which makes up a new Pauli operator), so we get a circuit with r Pauli rotations.

Lemma A2. Let

$$H = \sum_l p_l P_l, \quad (\text{A30})$$

where $P_l^2 = I$ and the p_l are a probability distribution. Then the random variable Z resulting from Hadamard tests on RTE with r steps and step size τ has expectation value

$$\mathbb{E}BZ(\tau, r) = g(\tau r) = \langle \psi | \exp(-i\tau r H) | \psi \rangle. \quad (\text{A31})$$

Here $B \leq \exp(\tau^2 r)$.

We now briefly comment on the difference between RTE and qDRIFT. First of all, if one chooses τ and r with $r = \Omega(\tau^2)$, they are very similar. At every step, the probability of sampling order $n \geq 2$ in RTE scales is $\mathcal{O}(\tau^2)$, so the expected number of steps where you sample $n \geq 2$ is $\mathcal{O}(r\tau^2) = \mathcal{O}(1)$ out of $r = \mathcal{O}(\tau^2)$ steps. The main difference is that RTE, with a Hadamard test, gives an *unbiased* estimator of the signal $g(t)$. This is convenient for analysis. On the other hand, we have also computed a simple expression for the exact signal resulting from qDRIFT, from which one can extract the same information about the energies. For qDRIFT, the damping depends on the energy as well; if E_0 is close to its minimal value of -1 (so the Hamiltonian is close to frustration free) this means that the damping factor is smaller.

3. Partially randomized product formulas

In Sec. V we have discussed a partially randomized scheme, which treats certain terms in the Hamiltonian deterministically, and the remainder randomly, and here we explain and analyze this scheme in more detail. The error analysis is particularly simple, where the signal from the partially randomized product formula has a bias from the deterministic decomposition, and a decreased amplitude due to the randomization. This intuition is illustrated in Fig. 11. We now recall the setup. We write

$$H = \underbrace{\sum_{l=1}^{L_D} H_l}_{=H_D} + \underbrace{\sum_{m=1}^M h_m P_m}_{=H_R}, \quad \lambda_R = \sum_m |h_m|. \quad (\text{A32})$$

Here, the P_m should square to identity, but the H_l need not. As before, we let

$$H_R = \lambda_R \sum_m p_m P_m, \quad (\text{A33})$$

where the p_m form a probability distribution (and we have absorbed signs into the P_m).

We let $S_p(\delta)$ be a deterministic product formula for H , written as a sum of the H_l and treating H_R as a single term. This means we can write

$$S_p(\delta) = U_I \cdots U_1, \quad (\text{A34})$$

where each U_i is either of the form $\exp(-\delta_i H_l)$ for some real δ_i , or $U_i = \exp(-\delta_i H_R)$. Here the number of terms is $I \leq (L_D + 1)N_{\text{stage}}$, where N_{stage} is the number of stages of the product formula. For the standard Trotter formula of order p , we have $N_{\text{stage}} = 2 \times 5^{k-1}$. For the second order




FIG. 11. Intuition for the error analysis: deterministic product formulas give a discrete approximation which has a bias in the frequency, randomized product formulas give a damping of the amplitude (and no bias when using RTE); the partially randomized method has both.

Suzuki-Trotter formula, we get

$$S_2(\delta) = e^{-\frac{1}{2}i\delta H_1} \cdots e^{-\frac{1}{2}i\delta H_{L_D}} e^{-i\delta H_R} e^{-\frac{1}{2}i\delta H_{L_D}} \cdots e^{-\frac{1}{2}i\delta H_1}. \quad (\text{A35})$$

Let $U_{i_1}, \dots, U_{i_N}$ for $N \leq N_{\text{stage}}$ denote the unitaries where we time evolve along H_R , for times $\delta_{i_1}, \dots, \delta_{i_N}$. We have $\delta_{i_1} + \dots + \delta_{i_N} = \delta$, and $\tilde{\delta} := |\delta_{i_1}| + \dots + |\delta_{i_N}| = \mathcal{O}(\delta)$ (higher-order methods require evolving backward in time as well). Next, we will apply RTE to each of the U_{i_p} . To this end, we choose a number of steps r_{i_p} to simulate $\exp(-\delta_{i_p} H_R)$. We may write

$$U_{i_p} = \exp(-i\delta_{i_p} H_R) = B_{i_p} \sum_k q_{i_p,k} W_{i_p,k}, \quad (\text{A36})$$

$$B_{i_p} \leq \exp(\delta_{i_p}^2 \lambda_R^2 / r_{i_p})$$

as in Sec. A2, where each $W_{i_p,k}$ is a product of Paulis P_m and r_{i_p} Pauli rotations. Expanding each U_{i_p} in this fashion, and taking the s th power of the result, gives a linear combination of unitaries,

$$S_p(\delta)^s = B \sum_j q_j W_j, \quad (\text{A37})$$

where

$$B = (B_{i_1} \cdots B_{i_N})^s \leq \exp\left(s \sum_p \delta_{i_p}^2 \lambda_R^2 / r_{i_p}\right). \quad (\text{A38})$$

Each W_j consists of $s(r_{i_1} + \dots + r_{i_N})$ Pauli rotations from the expansion of the evolution along H_R , and at most

$N_{\text{stage}}L_{DS}$ applications of $\exp(-\delta H_I)$ for some values of δ . Now, given δ and s , in order to keep the sample overhead constant, we choose $r_{ip} = \lceil \lambda_R^2 |\delta_{ip}| \delta s \rceil$, which gives

$$s \sum_p \delta_{ip}^2 \lambda_R^2 / r_{ip} \leq 1. \quad (\text{A39})$$

The total number of Pauli rotations applied in the randomized part of the resulting circuit is

$$\begin{aligned} r &= s \sum_p r_{ip} = s \sum_p \lceil \lambda_R^2 \delta_{ip} \delta s \rceil \leq \lambda_R^2 \tilde{\delta}^2 s^2 + s N_{\text{stage}} \\ &= \mathcal{O}(\lambda_R^2 \delta^2 s^2). \end{aligned} \quad (\text{A40})$$

The term $s N_{\text{stage}}$ comes from having to round the values of the r_{ip} to integers and is a small subleading contribution. We summarize this result as follows.

Lemma A3. Let H be given as in Eq. (A32) and let $S_p(\delta)$ be a deterministic product formula for H , written as a sum of the H_I and treating H_R as a single term. Then the random variables $\mathbf{Z}$ which results from the measurement outcomes of a Hadamard test, using the partially randomized product formula described above, has expectation value

$$\mathbb{E} \mathbf{BZ}(\delta, s) = \langle \psi | S_p(\delta)^s | \psi \rangle. \quad (\text{A41})$$

Here $B = \mathcal{O}(1)$, and the circuit consists of at most $N_{\text{stage}}L_{DS}$ time evolution unitaries along one of the H_I , and $r = \mathcal{O}(\lambda_R^2 \delta^2 s^2)$ time evolution unitaries along one of the P_m .

We note that the second-order Trotter-Suzuki formula has the benefit in this context that it only evolves in the positive time direction, so we have $\tilde{\delta} = \delta$. The conclusion of Lemma A3 is that from the Hadamard test with the partially randomized product formula, we get a signal which, up to a damping, implements the deterministic product formula which treats H_R as a single term. The resulting Trotter error can be bounded using any standard technique (such as commutator bounds).

Different schemes for composite random and deterministic product formulas have been proposed [14,16,75], treating a relatively small number of important terms using deterministic Trotterization and a tail of terms with small weight randomly. The scheme described above is the same as in Ref. [16], where we have replaced qDRIFT by RTE. Compared to the diamond norm bounds of Ref. [16], the randomized part of the product formula does not introduce any error. Additionally, in our analysis it is clear that if we are only interested in ground-state energies, $S_p(\delta)$ only needs to be accurate on the ground state, a significantly weaker requirement on the Trotter error. This allows for a sharper analysis of the performance of

partially randomized product formulas for single-ancilla phase estimation.

A final comment is that a general potential disadvantage of (partially) randomized product formulas is that in principle for each Hadamard test one has to sample a new random circuit. In practice, it may be more convenient (e.g., for compilation purposes) to repeat the same circuit multiple times. For randomized product formulas [51] as well as partially randomized product formulas [27] it has been shown that this is in fact possible. These works show that with increasing circuit depth, the expectation value of observables (such as Hadamard test outcomes) rapidly concentrate with respect to the randomness of the sampled circuit.

4. Halving the required time evolution

There is a basic trick to reduce the maximum required evolution time when using the Hadamard test for time evolution [17]. This assumes we have a symmetric Trotter formula (meaning that $S_p(-t) = S_p(t)^\dagger$), which is the case for the even order Suzuki-Trotter formulas. We write out the Trotter unitary in terms of the individual Hamiltonian time evolutions

$$S_p(\delta) = U_I(\delta_I) \cdots U_1(\delta_1), \quad (\text{A42})$$

where each unitary is time evolution along a local term H_I of the Hamiltonian. Now, for some arbitrary unitary U , let $C_\pm U$ denote a controlled version

$$C_\pm U = |0\rangle\langle 0| \otimes U^\dagger + |1\rangle\langle 1| \otimes U, \quad (\text{A43})$$

so the product formula, we have is

$$C_\pm U_i(\delta_i) = |0\rangle\langle 0| \otimes e^{i\delta_i H_{I_i}} + |1\rangle\langle 1| \otimes e^{-i\delta_i H_{I_i}}. \quad (\text{A44})$$

By the symmetry property of the product formula, we have

$$C_\pm S_p(\delta) = C_\pm U_I(\delta_I) \cdots C_\pm U_1(\delta_1). \quad (\text{A45})$$

It is now easy to see that applying the Hadamard test to the unitary $(C_\pm S_p(\delta))^s$ gives expectation value $\langle \psi | S_p(\delta)^{2s} | \psi \rangle$. This is implemented by the circuit in Fig. 12. If the $U_i(\delta_i)$ are Pauli rotations, then it requires only Cliffords and one single-qubit rotation to implement $C_\pm U_i(\delta_i)$. Therefore, this trick reduces the required number of single-qubit rotations by a factor of 2.

This trick does *not* apply for randomized product formulas (which do not have the right symmetry property). However, it can be modified to apply for partially randomized product formulas. In that case, one proceeds as above, thinking of H as consisting of $H_1, \dots, H_{L_D}$ and H_R as in Eq. (A32). This gives a circuit with controlled applications of unitaries of the form $U = \exp(\pm i\delta H_R)$. Finally,




FIG. 12. Performing the Hadamard test on $C_{\pm S_p}(\delta)^s$ allows one to estimate $\langle \psi | S_p(\delta)^{2s} | \psi \rangle$. If $U = e^{i\delta P}$ is a Pauli rotation, we can let C be a Clifford which is such that $CPC^\dagger$ acts as a Pauli Z on the first qubit and use this to implement $C_{\pm}U$ with only one single-qubit rotation.

we expand

$$C_{\pm}U = (|0\rangle\langle 0| \otimes U^\dagger + |1\rangle\langle 1| \otimes I)(|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes U) \quad (\text{A46})$$

and separately apply RTE to U and $U^\dagger$ (with independent randomness). Altogether, this halves the gate count for the deterministic part of the partially randomized product formula (but not for the randomized part).

APPENDIX B: ROBUST PHASE ESTIMATION

As a phase estimation algorithm we consider *robust phase estimation* [35,41–43]. This algorithm is particularly simple and requires shorter depth if the guiding state is close to the ground state. Its main disadvantage is that it requires the ground-state overlap η to be at least some constant.

The idea of robust phase estimation is to estimate the time evolution signal at times $t = 2^m$ for $m = 0, 1, \dots, M$ with $M = \lceil \log \varepsilon^{-1} \rceil$. For each m the outcome of the Hadamard test $Z(2^m)$ has expectation value

$$g(2^m) = \sum_k c_k \exp(-i2^m E_k). \quad (\text{B1})$$

We repeat N_m times and let $\bar{Z}(2^m)$ denote the average. Given some outcome we compute an angle $\phi_m = \arg(\bar{Z}(2^m))$. Each m corresponds to one bit of precision for the estimate of E_0 , i.e., ϕ_m is an approximation of $2^m E_0$ modulo 2π . The set of estimates of E_0 compatible with ϕ_m is $\{2^{-m}(2\pi k + \phi_m) : k = 0, \dots, 2^m - 1\}$. Given a guess

θ_{m-1} for E_0 after $m - 1$ rounds, we let θ_m be given by

$$\theta_m = 2^{-m}(2\pi k + \phi_m), \quad (\text{B2})$$

for the integer $k = 0, \dots, 2^m - 1$ such that the angles θ_{m-1} and θ_m are as close as possible. For two angles θ, ϕ we let $d(\theta, \phi)$ denote the distance between the angles, so

$$d(\theta, \phi) = \min_{k \in \mathbb{Z}} |\theta - \phi + 2k\pi|. \quad (\text{B3})$$

The robust phase estimation algorithm is prescribed in detail in Algorithm 1. There, the angles ϕ_j are treated as input (and can, for example, be obtained by sampling Hadamard test measurement outcomes).

1. Rigorous bounds for robust phase estimation

We will now review the argument for the scaling of robust phase estimation and argue for what happens when using randomized simulation methods. The algorithm gives an accurate estimate of a target E if the estimate ϕ_m in round m is within distance at most $\pi/3$ of $2^m E$. If

ALGORITHM 1. Robust phase estimation.

Input: A sequence of angles ϕ_m , $m = 0, 1, \dots, M$.
Output: A phase estimate θ .

```

1  $\theta_{-1} \leftarrow 0$ 
2 for  $m = 0, 1, \dots, M$  do
3    $S_m = \{2^{-m}(2\pi k + \phi_m) : k = 0, \dots, 2^m - 1\}$ 
4    $\theta_m \leftarrow \arg \min_{\theta \in S_m} d(\theta, \theta_{m-1}), -\pi < \theta_m \leq \pi$ 
5 end for
6  $\theta \leftarrow \theta_M$ 
```

this is not the case, we will say that there is an error in round m . The Heisenberg scaling of robust phase estimation is achieved by making sure that the probability of error in early rounds is sufficiently small. The following results are derived in Refs. [35,41,42] and are the key to proving Heisenberg-limited scaling as well as robustness of the phase estimation algorithm.

Lemma B1. Let $\xi \leq 1$.

- (1) If in Algorithm 1, $d(\phi_m, 2^m E) < \pi/3$ for $m = 0, 1, \dots, m$, then θ_m in Algorithm 1 is such that $d(\theta_m, E) \leq 2^{-m}(\pi/3)$. If $d(\phi_m, 2^m E) < \pi/3$ for $m = 0, \dots, M-1$ and $d(\phi_M, 2^M E) \leq \xi$, then $d(\theta, E) \leq 2^{-M}\xi$.
- (2) Suppose that $\{\phi_m\}_{m=0}^M$ are independent real random variables, such that for $m = 1, \dots, M$ we have

$$\mathbb{P}(d(\phi_m, 2^m E) \geq \frac{\pi}{3}) \leq \xi^2 4^{-\alpha(M-m)}, \quad (\text{B4})$$

for some $\alpha > 1$, and

$$\mathbb{E} d(\phi_M, 2^M E)^2 \leq \xi^2. \quad (\text{B5})$$

Then applying Algorithm 1 to $\{\phi_m\}_{m=0}^M$ gives an estimate θ of E with

$$\mathbb{E} d(\theta, E)^2 \leq C \xi^2 4^{-M}. \quad (\text{B6})$$

Proof. The argument for the first claim is by induction. For $m = 0$, the result holds by assumption. Suppose that $d(\phi_m, 2^m E) < \pi/3$ and $d(\theta_{m-1}, E) \leq 2^{-m+1}(\pi/3)$. Then the candidate set S_m contains an estimate which is $(2^{-m}(\pi/3))$ -close to E given by $2^{-m}(2\pi k + \phi_m)$. Let k' be the integer such that $\theta_m = 2^{-m}(2\pi k' + \phi_m)$, then we would like to show that $k = k'$. The distance between candidates in S_m is at least $2^{-m+1}\pi$, so it suffices to show that $d(\theta_{m-1}, 2^{-m}(2\pi k + \phi_m)) < 2^{-m}\pi$. We use the triangle inequality, and the induction hypothesis that $d(\theta_{m-1}, E) \leq 2^{-m+1}(\pi/3)$ to confirm this:

$$\begin{aligned} d(\theta_{m-1}, 2^{-m}(2\pi k + \phi_m)) &\leq d(\theta_{m-1}, E) + d(E, 2^{-m}(2\pi k + \phi_m)) \\ &< 2^{-m+1}\frac{\pi}{3} + 2^{-m}\frac{\pi}{3} = 2^{-m}\pi. \end{aligned} \quad (\text{B7})$$

In the last round, if $d(\phi_M, 2^M E) \leq \xi$, then $d(\theta, E) \leq 2^{-M}\xi$. Next, if for $l \leq m$ we have $d(\phi_l, 2^l E) < \pi/3$ (but possibly not for $l > m$), then the output $\theta = \theta_M$ satisfies $d(\theta_M, E) \leq 2^{-m+3}(\pi/3)$. This follows from the observation that the algorithm is such that $d(\theta_l, \theta_{l+1}) \leq 2^{-l}\pi$,

so

$$\begin{aligned} d(\theta_M, E) &\leq d(\theta_m, E) + \sum_{l=m}^{M-1} d(\theta_l, \theta_{l+1}) \leq 2^{-m}\frac{\pi}{3} \\ &+ \underbrace{\sum_{l=m}^{M-1} 2^{-l}\pi}_{\leq 2^{-m+1}\pi} \leq 2^{-m+3}\frac{\pi}{3}. \end{aligned} \quad (\text{B8})$$

We now consider the case where the angles ϕ_m are random variables. Let p_m denote the probability that $m = 0, \dots, M-1$ is the smallest value for which we have an error (so $d(\phi_m, 2^m E) \geq \pi/3$), and let p_M denote the probability there are no errors for $m = 0, \dots, M-1$. Then we can bound the mean square error

$$\begin{aligned} \mathbb{E} d(\theta_M, E)^2 &\leq \sum_{m=0}^{M-1} p_m \left(\frac{8\pi}{3 \cdot 2^{m-1}} \right)^2 \\ &+ p_M 4^{-M} \mathbb{E} d(\phi_M, 2^M E)^2. \end{aligned} \quad (\text{B9})$$

We may bound p_m by the probability of having an error in round m , so $p_m \leq \xi^2 4^{-\alpha(M-m)}$, and $p_M \leq 1$ which gives

$$\mathbb{E} d(\theta_M, E)^2 \leq \sum_{m=0}^{M-1} \xi^2 4^{-\alpha(M-m)} \left(\frac{8\pi}{3 \cdot 2^{m-1}} \right)^2 + \xi^2 4^{-M} \quad (\text{B10})$$

$$= \left(1 + \left(\frac{16\pi}{3} \right)^2 \sum_{m=0}^{M-1} 4^{-(\alpha-1)(M-m)} \right) \xi^2 4^{-M} \quad (\text{B11})$$

$$\begin{aligned} &\leq \left(1 + \frac{1}{4^{\alpha-1}-1} \left(\frac{16\pi}{3} \right)^2 \right) \xi^2 4^{-M} \\ &= C \xi^2 4^{-M} \end{aligned} \quad (\text{B12})$$

for a constant C which only depends on α . $\blacksquare$

In Appendix B2 we will give numerical estimates for the relevant constant factors. The angles ϕ_m are obtained from a Hadamard test, simulating time evolution for time 2^m . We will assume that we have a guiding state $|\psi\rangle$ with sufficiently high ground-state overlap. We first repeat the following fact from Ref. [43]:

Suppose $Z \in \mathbb{C}$, $d_k \geq 0$ and $E_k \in \mathbb{R}$. If

$$\left| Z - \sum_k d_k \exp(iE_k) \right| \leq \beta, \quad \frac{\beta + \sum_{k \geq 1} d_k}{d_0} \leq 1, \quad (\text{B13})$$

then we can bound

$$|Z - d_0 \exp(iE_0)| \leq \beta + \sum_{k \geq 1} d_k, \quad (\text{B14})$$

which implies that the sine of the angle between Z and $d_0 \exp(iE_0)$ satisfies $|\sin(d(E_0, \arg(Z)))| \leq (\beta + \sum_{k \geq 1} d_k)/d_0$ and hence

$$d(E_0, \arg(Z)) \leq \arcsin\left(\frac{\beta + \sum_{k \geq 1} d_k}{d_0}\right). \quad (\text{B15})$$

We run robust phase estimation on a signal estimated by a Hadamard test and estimate the energy E_0 based on the outcome $\theta = \theta_M$. We consider two variants. For the randomized methods, we assume we have a normalized Hamiltonian $H = \sum_l p_l P_l$ (in general that means we have to rescale ε with λ).

- (i) We estimate ϕ_j from applying a Hadamard test to e^{-iHt} , i.e., from the exact time evolution, for time $t = 2^m$. We use $2N_m$ samples to get an estimate Z_m of $g(2^m) = \langle \psi | e^{-iH2^m} | \psi \rangle$, and let $\phi_m = -\arg(Z_m)$. We estimate E_0 by θ_M . The total amount of time evolution we use is

$$t_{\text{tot}} = \sum_{m=0}^M 2N_m 2^m. \quad (\text{B16})$$

- (ii) We estimate ϕ_j from applying Hadamard tests using a randomized product formula (either qDRIFT or RTE) for time evolution $t = 2^m$ with $r_m = \Omega(2^{2m})$ Pauli rotations, and time step $\tau_m = t/r_m$. We use $2N_m$ samples to get an estimate Z_m of the signal, and let $\phi_m = -\arg(Z_m)$. For RTE, we estimate E_0 by θ_M , and in the case of qDRIFT, we estimate E_0 by $\tau_M^{-1} \tan(\tau_M \theta_M)$. In both cases, the total number of Pauli rotations used is

$$r_{\text{tot}} = \sum_{m=0}^M 2N_m r_m. \quad (\text{B17})$$

The argument for the exact time evolution model in the following result follows [35,41,42] for (1) and [43] for (3). A difference compared to Ref. [43] is that for (2) we bound the mean square error, instead of bounding a success probability for an error guarantee. The application to randomized product formulas in (2) and (4) is a straightforward consequence of the analysis.

Theorem B1. Suppose that we have a guiding state $|\psi\rangle$ with $c_0 \geq \eta > 4 - 2\sqrt{3} \approx 0.53$.

- (1) In the exact time evolution model, we obtain an estimate with root-mean-square error ε using $t_{\text{max}} = \mathcal{O}(1/\varepsilon)$, and $t_{\text{tot}} = \mathcal{O}(1/\varepsilon)$.
 (2) When using RTE or qDRIFT, we obtain an estimate with root-mean-square error ε using maximal number of rotations per circuit $r_{\text{max}} = \mathcal{O}(1/\varepsilon^2)$ and total number of rotations $r_{\text{tot}} = \mathcal{O}(1/\varepsilon^2)$.

In the limit $\eta \rightarrow 1$, for $\xi > \arcsin((1 - \eta)/\eta)$, we can reduce the maximal depth.

- (3) In the exact time evolution model, we obtain an estimate with root-mean-square error ε using $t_{\text{max}} = \mathcal{O}(\xi/\varepsilon)$, and $t_{\text{tot}} = \mathcal{O}(1/\xi\varepsilon)$.
 (4) When using RTE or qDRIFT, we obtain an estimate with root-mean-square error ε using maximal number of rotations per circuit $r_{\text{max}} = \mathcal{O}(\xi^2/\varepsilon^2)$ and total number of rotations $r_{\text{tot}} = \mathcal{O}(1/\varepsilon^2)$.

Proof. We estimate X_m and Y_m using N_m Hadamard tests and, thus, estimate $Z_m = X_m + iY_m$ using $2N_m$ Hadamard tests. For the case where we use a randomized product formula, we take $r_m = 2^{m+M}$. The expectation value of Z_m is given by

$$\mathbb{E}Z_m = \sum_k d_k e^{-i2^m F_k}, \quad (\text{B18})$$

where $F_k = E_k$ for the case of exact time evolution and RTE, and $F_k = 2^M \arctan(2^{-M} E_k)$ for qDRIFT. Note that if $|F_0 - \theta_M| \leq \varepsilon$, then

$$|E_0 - 2^M \tan(2^{-M} \theta_M)| \leq 2^M |\tan(2^{-M} \theta_M) - \tan(2^{-M} F_0)| \quad (\text{B19})$$

$$\leq \varepsilon + \mathcal{O}(2^{-2M} \varepsilon), \quad (\text{B20})$$

so it suffices to get an accurate estimate for F_0 . For qDRIFT, we assume for convenience that $|E_0| \geq |E_k|$ for $k \geq 1$ (which can always be achieved by shifting H with a constant factor).

For exact time evolution, $d_k = c_k$, for the randomized methods $e^{-1} c_k \leq d_k \leq c_k$ (while d_k also depends on m , this bound holds for all m). In particular, $d_0 \geq e^{-1} c_0 \geq e^{-1} \eta$. We let $0 < \beta = \eta(1 + \sin(\pi/3)) - 1$ for the case with exact time evolution and $0 < \beta = e^{-\frac{1}{2}}(\eta(1 + \sin(\pi/3)) - 1)$ for the randomized cases. This choice of β is such that Eq. (B13) is satisfied

$$\arcsin\left(\frac{\beta + \sum_{k \geq 1} d_k}{d_0}\right) \leq \frac{\pi}{3}. \quad (\text{B21})$$

Now, by Eq. (B15) and Hoeffding's inequality we see that for $\phi_m = -\arg(Z_m)$,

$$\begin{aligned} \mathbb{P}\left(d(\phi_m, F_0) \geq \frac{\pi}{3}\right) &\leq \mathbb{P}(|Z_m - \mathbb{E}Z_m| \geq \beta) \\ &\leq \mathbb{P}\left(|X_m - \mathbb{E}X_m| \geq \frac{\beta}{\sqrt{2}}\right) + \mathbb{P}\left(|Y_m - \mathbb{E}Y_m| \geq \frac{\beta}{\sqrt{2}}\right) \\ &\leq 4 \exp\left(-\frac{N_m \beta^2}{4}\right). \end{aligned} \quad (\text{B22})$$

In particular, we find that by choosing $N_m = \beta^{-2}(2 \log(\xi^{-1}) + \mathcal{O}(M - m))$ we can achieve

$$\mathbb{P}(d(\phi_m, 2^m F_0) \geq \frac{\pi}{3}) \leq \xi^2 4^{-\alpha(M-m)}, \quad (\text{B23})$$

which is a necessary condition for the RPE algorithm to succeed (see Lemma B1). We now discuss the first part of the theorem and take ξ constant. By Lemma B1 and Eq. (B23) ε^2 is bounded as $\mathcal{O}(4^{-M})$. The total required evolution time is given by

$$t_{\text{tot}} = \sum_{m=0}^M 2N_m 2^m = \mathcal{O}(2^M) \quad (\text{B24})$$

and similarly $r_{\text{tot}} = \mathcal{O}(2^{2M})$. This results in $t_{\text{tot}} = \mathcal{O}(\varepsilon^{-1})$ and $r_{\text{tot}} = \mathcal{O}(\varepsilon^{-2})$ and proves statements (1) and (2) of the theorem.

For (3) and (4), we choose $N_M = \mathcal{O}(\beta^{-2}\xi^{-2})$ samples in the last round. To bound the error, we bound by

$$\begin{aligned} \mathbb{E} d(\phi_M, 2^M F_0)^2 &\leq \mathbb{P}(|Z_M - \mathbb{E}Z_M| \geq \beta) \pi^2 \\ &+ \mathbb{E} \left[d(\phi_M, 2^M F_0)^2 \mid |Z_M - \mathbb{E}Z_M| < \beta \right]. \end{aligned} \quad (\text{B25})$$

The first term is bounded by Eq. (B22), so $\mathbb{P}(|Z_M - \mathbb{E}Z_M| \geq \beta) = 4 \exp(-\Omega(\xi^{-2})) = \mathcal{O}(\xi^{-2})$. For the second term, we use that for $\delta = |Z_M - \mathbb{E}Z_M| < \beta$ we have

$$d(\phi_M, 2^M F_0) \leq \arcsin \left(\frac{\delta + (1 - \eta)}{\eta} \right) = \mathcal{O}(\delta + (1 - \eta)), \quad (\text{B26})$$

(which for Lemma B1 we would like to be $\mathcal{O}(\xi^2)$). Here we use Eq. (B15), $\eta^{-1}(\delta + (1 - \eta)) \leq \pi/3$ and η is close to 1. This can be used to see that the second term in Eq. (B25) can be bounded by

$$\begin{aligned} \mathbb{E} \arcsin \left(\frac{|Z_M - \mathbb{E}Z_M| + (1 - \eta)}{\eta} \right)^2 \\ = \mathcal{O} \left(\mathbb{E} |Z_M - \mathbb{E}Z_M|^2 + (1 - \eta)^2 \right) = \mathcal{O}(\xi^2), \end{aligned} \quad (\text{B27})$$

if $N_M = \Omega(\xi^{-2})$. By Lemma B1 and Eq. (B23), this implies root-mean-square error $\varepsilon = \mathcal{O}(\xi 2^{-M})$. The maximal evolution time is $t_{\text{max}} = 2^M = \mathcal{O}(\xi \varepsilon^{-1})$, and the randomized methods use $r_{\text{max}} = r_M = \mathcal{O}(\xi^2 \varepsilon^{-2})$, and

$$\begin{aligned} t_{\text{tot}} &= \sum_{m=0}^M 2N_m 2^m = \sum_{m=0}^M \mathcal{O}(\beta^{-2}(\log(\xi^{-1}) + (M - m))2^m) \\ &+ \mathcal{O}(\xi^{-2}2^M) = \mathcal{O}(\xi^{-2}2^M) = \mathcal{O}(\xi^{-1}\varepsilon^{-1}). \end{aligned} \quad (\text{B28})$$

With the randomized methods, the total number of rotations is similarly given by

$$r_{\text{tot}} = \sum_{m=0}^M 2N_m r_m = \mathcal{O}(\xi^{-2}2^{2M}) = \mathcal{O}(\varepsilon^{-2}). \quad (\text{B29})$$

■

We briefly comment on the choice of depth r_m for the randomized product formulas, what choice minimizes the total cost, and what is the resulting difference between RTE and qDRIFT. Recall that the depth chosen determines how much the amplitude of the signal is damped, as in Eqs. (A15) and (A29). If the relevant component of the signal is damped by a factor B , we need to scale the number of samples by $\mathcal{O}(B^2)$. For RTE for time t and r rotations, we get a damping factor $B \leq \exp(t^2/r)$, which means that we have to take a number of samples scaling with $B^2 \leq \exp(2t^2/r)$ to estimate the signal to constant precision. The total cost scales as $B^2 r$, which is optimized by $r = 2t^2$. In particular, for $t = 2^m$, we should choose $r_m = 2^{2m+1}$. For qDRIFT, for our proof analysis above it was convenient to choose $r_m = 2^{M+m}$, since that means that we have a constant value of F_k over different rounds m . Alternatively, one may take $r_m = 2^{2m}$, in which case in round m we have $2^m F_0 = 2^{2m} \arctan(2^{-m} E_0) = 2^m E_0 + \mathcal{O}(2^{-m})$. Since we only need to estimate $2^m E_0$ to constant precision in the m th round, this does not make a significant difference (one has to be careful for small m , but in that case it is anyways cheap to take r_m larger). By Eq. (A15), we have $c_0 \geq \exp(-t^2(1 - E_0^2)/2r) \geq \exp(-t^2/2r)$. In this case, optimizing the total resources means one should take $r = t^2$ (so this is a factor of 2 more efficient than RTE). One caveat is that the damping depends on E_k for qDRIFT, which makes a difference if we do not start with the ground state. This is beneficial if $|E_0| \leq |E_k|$ for all $k \neq 0$ with $c_k > 0$ (the excited states are damped more than the ground state).

2. Performance of robust phase estimation

In the previous section we saw that robust phase estimation achieves Heisenberg-limited scaling, when choosing an appropriate number of samples for the Hadamard test. There, we did not attempt to optimize constant factors. In the special case where we start with an exact eigenstate, Ref. [42] shows a bound for the root-mean-square error of $\varepsilon \leq C_{\text{tot}} \pi / t_{\text{tot}}$ with $C_{\text{tot}} \approx 25$ using $N_m = \lceil N_M + (M - m)D \rceil$ samples with $N_M = 11$ and $D \approx 4.11$, based on the analysis in Lemma B1. This is in contrast to the optimal performance of QPE in general, which has optimal constant factor $C_{\text{tot}} = 1$ [76]. However, the rigorous bound of Ref. [42] is still loose, and for our resource estimates we use an empirically estimated constant. Numerical simulation shows that using the choice

of parameters from Ref. [42], we achieve a scaling with a constant factor of $C_{\text{tot}} \approx 5$, see Fig. 13(a). The empirical scaling with the *maximal* time evolution $t_{\text{max}} = 2^M$ is $\varepsilon \approx C_{\text{max}}\pi/t_{\text{max}}$ for $C_{\text{max}} \approx 0.08$. These results corroborate previous numerical estimates, see e.g., Ref. [77]. The total time evolution is given by

$$t_{\text{tot}} = \sum_{m=0}^M 2N_m 2^m = \sum_{m=0}^M 2(N_M + D(M-m))2^m \leq 4(N_M + D)2^M = 4(N_M + D)t_{\text{max}}. \quad (\text{B30})$$

This estimate only ignores a small subleading term, and gives a result which is consistent with the numerical relation between C_{tot} and C_{max} from Fig. 13. These estimates assume a perfectly accurate quantum computer. By design, the robust phase estimation procedure can tolerate a small amount of error. See Ref. [77] for a discussion of the error budget in the context of quantum chemistry computations.

We now discuss the total cost when using randomized product formulas. For RTE, using $\tau = 2^{-m-1}$ and $r = 2^{2m+1}$ in the m th round requires a sampling overhead of e . We should now take N_m a factor e larger than in the deterministic case. For our numerical estimates, we take $N_m = e(11 + 4(M-m))$. The total number of rotations is given to good approximation by

$$r_{\text{tot}} = \sum_{m=0}^M 2N_m 2^{2m+1} = \sum_{m=0}^M 2(N_M + D(M-m))2^{2m+1} \leq 16 \left(\frac{N_M}{3} + \frac{D}{9} \right) 2^{2M} = 16 \left(\frac{N_M}{3} + \frac{D}{9} \right) t_{\text{max}}^2 = 8 \left(\frac{N_M}{3} + \frac{D}{9} \right) r_{\text{max}}. \quad (\text{B31})$$

In Fig. 13(b) we numerically show that this indeed leads to a scaling $\varepsilon^2 \approx C'_{\text{tot}}/r_{\text{tot}}$ and $\varepsilon^2 \approx (C'_{\text{max}}\pi)^2/r_{\text{max}}$ where $C'_{\text{tot}} \approx 16$ and $C'_{\text{max}} \approx 0.2$ (note that the maximal number of rotations is $2r_{\text{max}}^2$). As discussed at the end of the previous subsection, for qDRIFT we save an additional factor of 2 for both r_{max} and r_{tot} , giving a total cost of approximately $r_{\text{tot}} = 8\varepsilon^{-2}$ for performing phase estimation with root-mean-square error ε .

Concretely, for the example of FeMoco discussed in the introduction, we have $\lambda = 405$ and $\varepsilon = 0.0016$. We should estimate the ground-state energy of the normalized Hamiltonian to precision $\lambda^{-1}\varepsilon$. We choose $\xi = 0.1$, giving $M = \lceil \log_2(\xi\lambda\varepsilon^{-1}) \rceil = 15$. For the last round, we do not necessarily need to take a power of two, but we can choose $2^{M-1} < K_M \leq 2^M$ and estimate $K_M E_0$. We let $K_M = \lceil \xi\lambda\varepsilon^{-1} \rceil$, then estimating $K_M E_0$ to precision ξ gives an estimate of E_0 to precision ε . This gives circuits with at most $K_M^2 = 6.4 \times 10^8$ Pauli rotations.

APPENDIX C: THE ELECTRONIC STRUCTURE PROBLEM

The motivation of this work is to use quantum computing for Hamiltonians which arise in electronic structure problems in the Born-Oppenheimer approximation, meaning that we fix the location of nuclei, and treat the electron wavefunctions as the quantum degrees of freedom. For completeness, we describe the setup, relevant to the numerical examples, in some detail here. We work in second quantization and we assume that we have chosen a set of orthonormal spatial orbitals labeled by an index $p \in [N]$ corresponding to a function $\phi_p \in L^2(\mathbb{R}^3)$. Additionally, there is a spin degree of freedom, so we end up with a set of spin-orbitals $\phi_{p,\alpha}$ for $p \in [N]$ and $\alpha = \uparrow, \downarrow$. For convenience, we will assume that the spatial orbitals are real-valued functions (but this is not essential to the arguments). The spin-orbitals span a single-particle space $\mathcal{H}_1 \cong \mathbb{C}^{2N}$. The Hilbert space we work with is the corresponding Fock space

$$\mathcal{H} = \bigoplus_{n=0}^{2N} \mathcal{H}_n \quad \mathcal{H}_n = \mathcal{H}_1^{\wedge n}. \quad (\text{C1})$$

We denote the fermionic annihilation operator associated to orbital p and spin α by $a_{p\sigma}$. The size of the problem is determined by N , the number of orbitals, and n , the number of electrons.

Given a choice of orbitals, the Hamiltonian of the electronic structure problem is given by

$$H = T + V = \sum_{pq} \sum_{\sigma} h_{pq} a_{p\sigma}^\dagger a_{q\sigma} + \frac{1}{2} \sum_{pqrs} \sum_{\sigma\tau} h_{pqrs} a_{p\sigma}^\dagger a_{r\tau}^\dagger a_{s\tau} a_{q\sigma}, \quad (\text{C2})$$

where the coefficients h_{pq} and h_{pqrs} result from performing appropriate integrals involving the orbital functions. Here and in the remainder of this work we use the convention that the spatial orbitals are labeled by $p, q, r, s, \dots$ and spins by $\sigma, \tau, \dots$. In Eq. (C2), the first term T represents the kinetic energy and the interaction between the nuclei and the electrons

$$h_{pq} = \int_{\mathbb{R}^3} \phi_p(\vec{r}) \left(-\frac{1}{2} \nabla^2 - \sum_{i=1}^m \frac{Z_i}{\|\vec{r} - \vec{R}_i\|} \right) \phi_q(\vec{r}) d\vec{r}. \quad (\text{C3})$$

Here the $\vec{R}_i$ and Z_i are the (fixed) positions and the nuclear charge of the m nuclei. The last term corresponds to the




FIG. 13. Estimation of the overhead of robust phase estimation. Distribution of the mean square error, 10^4 data points for each value of M . The top panels show scaling of mean square error with total time evolution and maximal time evolution $t_{\text{max}} = 2^M$, using $N_m = 11 + 4(M - m)$ samples in round m . The bottom panels show scaling of mean square error with total number of rotations according to Eq. (B29) and maximal time evolution $t_{\text{max}} = 2^M$ simulated, using depth $r_m = 2^{2m+1}$ and $N_m = \lceil e(11 + 4(M - m)) \rceil$ samples in round m .

electron-electron Coulomb interaction V ,

$$h_{pqrs} = \int_{\mathbb{R}^3} \int_{\mathbb{R}^3} \frac{\phi_p(\vec{r}_1)\phi_q(\vec{r}_1)\phi_r(\vec{r}_2)\phi_s(\vec{r}_2)}{\|\vec{r}_1 - \vec{r}_2\|} d\vec{r}_1 d\vec{r}_2. \quad (\text{C4})$$

The integrals depend on the choice of the basis of orbitals. Any unitary $u \in U(N)$ defines a basis change on the spatial part of the single-particle space, giving rise to a set of orbitals by $\tilde{\phi}_p = \sum_q u_{pq} \phi_q$. It defines a particle-number and spin conserving unitary U on the Hilbert space $\mathcal{H}$ by defining a set of fermionic creation operators $\tilde{a}_p^\dagger = \sum_q u_{pq} a_q^\dagger$. Typically, the set of orbitals used to define Eq. (C2) is derived from a mean-field calculation, giving rise to a set of *canonical orbitals*. Often it is useful to perform a basis change so the orbitals are *localized* functions with exponentially decaying tails. In that case, from

Eq. (C4) we see that the density-density terms with $p = q$ and $r = s$ lead to contributions which decay polynomially in the distance between the spatial locations of the orbitals p and r , whereas for $p \neq q$ we have an exponential decay in the distance between the spatial locations of orbitals p and q . In Fig. 8 we indeed see a regime of terms with power law decay, and a regime of terms with exponential decay. For large spatially extended systems, one finds that many of the h_{pqrs} are extremely small (say, less than 10^{-10}) and can be safely ignored. In that case, the number of relevant terms scales as $\mathcal{O}(N^2)$. Details on electronic structure formalism and computations can be found in Ref. [78]; see in particular Sec. 9.12 for a detailed discussion of the scaling of the size of integrals h_{pqrs} .

The Hamiltonian H is mapped to $2N$ qubits by a fermion-to-qubit mapping. Various options are possible,

but the key feature is that the terms in H get mapped to Pauli operators. Generally, one defines the Majorana fermionic operators by

$$\gamma_{p\sigma,0} = a_{p\sigma}^\dagger + a_{p\sigma}, \quad \gamma_{p\sigma,1} = i(a_{p\sigma}^\dagger - a_{p\sigma}) \quad (\text{C5})$$

and each $\gamma_{p\sigma,\alpha}$ should be mapped to a Pauli operator in a way that preserves the anticommutation relation $\{\gamma_{p\sigma,\alpha}, \gamma_{q\tau,\beta}\} = 0$ if $p\sigma,\alpha \neq q\tau,\beta$ (for example, using the Jordan-Wigner mapping, or Bravyi-Kitaev mapping). Ignoring a component which is proportional to the identity operator, the Hamiltonian can be written as

$$H = \frac{i}{2} \sum_{pq} \sum_{\sigma} \left(h_{pq} + \sum_r h_{pqrr} - \frac{1}{2} \sum_r h_{prrq} \right) \gamma_{p\sigma,0} \gamma_{q\sigma,1} - \frac{1}{8} \sum_{pqrs} \sum_{\sigma\tau} h_{pqrs} \gamma_{p\sigma,0} \gamma_{q\sigma,1} \gamma_{r\tau,0} \gamma_{s\tau,1}. \quad (\text{C6})$$

A minimal representation in which each Majorana operator products appears at most once has been derived in Refs. [60,79], which gives a representation of H ,

$$H = \sum_{l=1}^L h_l P_l, \quad (\text{C7})$$

where the P_l are Pauli operators and the h_l are real coefficients.

1. Factorization of electronic structure Hamiltonians

The Coulomb interaction (i.e., the two-body term) is the part of the Hamiltonian that is the main contribution to the complexity of quantum simulation of H in Eq. (C2). For example, simply Trotterizing over the sum as given in the description of the Hamiltonian, the one-body interactions give $\mathcal{O}(N^2)$ terms, while the two-body interaction contributes $\mathcal{O}(N^4)$ terms. Moreover, if the Hamiltonian *only* has one-body terms, the Hamiltonian is quadratic in the creation and annihilation operators, which may be simulated efficiently classically. Employing anticommutation relations the electronic structure Hamiltonian C2 can be written as

$$H = \sum_{pq} \sum_{\sigma} \underbrace{\left(h_{pq} - \frac{1}{2} \sum_r h_{prrq} \right)}_{h'_{pq}} a_{p\sigma}^\dagger a_{q\sigma} + \frac{1}{2} \sum_{pqrs} \sum_{\sigma\tau} h_{pqrs} a_{p\sigma}^\dagger a_{q\sigma} a_{r\tau}^\dagger a_{s\tau} = T + V. \quad (\text{C8})$$

It is useful to look for simpler representations of the Coulomb interaction V in order to simplify both product formulas and block encodings of the Hamiltonian.

A prominent example is *double factorization* [2,18,19,24,80–82].

The idea is to first consider h as a positive $N^2 \times N^2$ real matrix (the positivity follows from Eq. (C4)) and decompose it as

$$h_{pqrs} = \sum_{j=1}^L (\mathcal{L}_{pq}^{(j)})^T \mathcal{L}_{rs}^{(j)}, \quad (\text{C9})$$

where $\mathcal{L}^{(j)}$ is a real vector of dimension N^2 , with entries labeled by orbital pairs pq . This is the “first factorization.” The Coulomb tensor satisfies the symmetries $h_{pqrs} = h_{pqsr} = h_{qprs} = h_{qpsr} = h_{rspq} = h_{rsqp} = h_{srpq} = h_{srqp}$. These symmetries imply that we may take the $\mathcal{L}^{(j)}$ to be real and such that $\mathcal{L}_{pq}^{(j)} = \mathcal{L}_{qp}^{(j)}$. Therefore, we can reinterpret the $\mathcal{L}^{(j)}$ as *symmetric* $N \times N$ matrices. One may now apply a “second factorization” and diagonalize each $\mathcal{L}^{(j)}$ using an orthogonal transformation. That is, we find $u^{(j)} \in O(N)$ and diagonal $N \times N$ matrices $\Lambda^{(j)}$ with real numbers $\lambda_k^{(j)}$ on the diagonal, such that

$$\mathcal{L}^{(j)} = (u^{(j)})^T \Lambda^{(j)} u^{(j)}. \quad (\text{C10})$$

Inserting this back into the Coulomb interaction term of the Hamiltonian, one finds

$$V = \frac{1}{2} \sum_{j=1}^L \left(\sum_{pq} \sum_{\sigma} \mathcal{L}_{pq}^{(j)} a_{p\sigma}^\dagger a_{q\sigma} \right)^2 = \frac{1}{2} \sum_{j=1}^L \left(\sum_{k=1}^{\rho_j} \lambda_k^{(j)} \sum_{pq} \sum_{\sigma} u_{kp}^{(j)} a_{p\sigma}^\dagger a_{q\sigma} u_{kq}^{(j)} \right)^2. \quad (\text{C11})$$

The $N \times N$ matrices $u^{(j)}$ define a change of orbital basis, and they give rise to a corresponding orbital rotation $U^{(j)}$ on the full Hilbert space. The fermionic annihilation operators in the changed basis $\phi_k^{(j)} = \sum_p u_{kp}^{(j)} \phi_p$ are given by

$$a_{k\sigma}^{(j)} = \sum_p u_{kp}^{(j)} a_{p\sigma} \quad (\text{C12})$$

and the associated number operators are $n_{k\sigma}^{(j)} = (a_{k\sigma}^{(j)})^\dagger a_{k\sigma}^{(j)}$. Then,

$$V = \frac{1}{2} \sum_{j=1}^L \left(\sum_{k=1}^{\rho_j} \sum_{\sigma} \lambda_k^{(j)} n_{k\sigma}^{(j)} \right)^2 = \sum_{j=1}^L (U^{(j)})^T V^{(j)} U^{(j)}, \quad \text{for } V^{(j)} = \frac{1}{2} \left(\sum_{k,\sigma} \lambda_k^{(j)} n_{k\sigma} \right)^2. \quad (\text{C13})$$

The main reason that this decomposition is useful is that h is well-approximated by a *low-rank* matrix. In the first factorization in Eq. (C9) we may truncate the sum to contain

only $L = \tilde{O}(N)$ terms, instead of the maximal number of N^2 . Secondly, numerical results suggest that in the limit of large N , for spatially extended systems, one can truncate ρ_j , the number of terms in the second factorization, to only $\tilde{O}(1)$ terms instead of N [65,66], although this only becomes relevant for large systems. Note that in this construction the choice of the $\mathcal{L}^{(j)}$ is not unique. Common choices are an eigenvalue decomposition or a Cholesky decomposition, but there is significant flexibility. There are various methods which aim to optimize the factorization in order to reduce the total circuit depth [19,81,83]. We discuss the associated gate counts for Trotterization using this low-rank factorization in Appendix E.

2. Benchmark systems

For FeMoco we use the model Hamiltonian first described by Reiher *et al.* [17]. There is an alternative widely used choice of active space given by Ref. [63] that is believed to capture the electronic structure more accurately, although even larger active spaces may be needed [84]. For our resource estimates we stick to the Hamiltonian from Ref. [17], for facilitating comparison with deterministic product formulas (which are only available from Ref. [17]). For the other methods, the active space from Ref. [63] shows a qualitatively similar picture. The data specifying the Hamiltonian coefficients for both the hydrogen chain and for FeMoco is taken from Ref. [4]. For the hydrogen chain we use a STO-6G minimal basis set and set the distance between the hydrogen atoms at 1.4 Bohr. The number of electrons is $n = N_H$, the number of spin-orbitals is $2N_H$. Additionally, we use a collection of small molecules with varying active space sizes to estimate the Trotter error. These molecules are alanine $C_3H_7NO_2$, the chromium dimer Cr_2 , an iron-sulfur cluster Fe_2S_2 , hydrogen peroxide H_2O_2 , imidazole $C_3H_4N_2$, naphthalene $C_{10}H_8$, nitrous acid HNO_2 and sulfurous acid H_2SO_3 at equilibrium geometries. The active space sizes vary from four to nine orbitals, and the active orbitals are chosen to be the closest ones to the highest occupied molecular orbital - lowest unoccupied molecular orbital (HOMO-LUMO) gap. Furthermore, for each active space Hamiltonian generated this way, the same λ optimization as described below is performed.

3. Optimizing the weight λ

As discussed in Sec. VII, it is useful to optimize the choice of representation to have small weight λ . Firstly, it is clear that the coefficients h_{pq} and h_{pqrs} defined by Eq. (C3) and Eq. (C4) depend on the choice of orbital basis. Typically, one starts from a set of canonical orbitals from a mean-field calculation, which determine the active space for the problem. However, given the subspace, one is free to choose a different basis, for which the representation of H has desirable properties, such as a small weight λ . In

Ref. [60] the effect of different basis changes on the value of λ was explored, with as conclusion that orbital localization (i.e., choosing an orbital basis for which the orbital functions are spatially localized) decreases the value of λ and that further reduction is possible by minimizing over the basis change. We perform such optimization for all benchmark systems, using gradient descent over the basis change unitaries to minimize the weight λ . The problem is not convex and needs a good starting point. One option is to use localized orbitals. From our numerics, we observe that this choice can be a local minimum for λ . A good alternative is to take the largest term in a Cholesky decomposition of the Coulomb tensor [cf. Eq. (C9)] $\mathcal{L}_{pq}^{(1)}$. This defines an $N \times N$ symmetric matrix, and we can consider the orbital basis which diagonalizes this matrix, as in Eq. (C10). This is cheap to compute, already leads to significantly reduced weight λ , and forms a good ansatz to start the optimization.

The Hamiltonian has a particle-number symmetry. That is, H commutes with the electron number operator $\hat{n} = \sum_{p,\sigma} a_{p\sigma}^\dagger a_{p\sigma}$, and we work in a subspace of fixed electron number n , which is set by the initial state. This feature can be exploited to define a different Hamiltonian, which equals H on the subspace of n -particles (possibly up to an additive constant factor) but with lower weight. This has been used in various ways, in particular to optimize the cost of LCU implementations of H [19,61,85]. For this work we use a variant of the method of Ref. [62]. Let $\hat{n} = \sum_{p,\sigma} n_{p\sigma}$ be the number operator, and let n be the number of electrons for the ground state of H . Then we can also apply phase estimation to $H + F(\hat{n} - n) + (\hat{n} - n)F$ for any self-adjoint operator F , since this has the same spectrum when restricted to the n -particle subspace. In particular, one can take F to be a one-body operator and use this to reduce the weight of the two-body terms in H . We perform a joint numerical optimization, minimizing λ , over the choice of the symmetry shift and the choice of orbital basis, using the limited-memory Broyden-Fletcher-Goldfarb-Shanno with bounds (L-BFGS-B) quasi-Newton optimization method. For the optimization, we restrict to an ansatz where we decompose the Coulomb operator as in Eq. (C13) and shift

$$V^{(j)} \rightarrow \frac{1}{2} \left(\sum_{k\sigma} \lambda_k^{(j)} n_{k\sigma} + f_j (\hat{n} - n) \right)^2 \quad (C14)$$

and optimize the parameters f_j . We show the resulting scaling of λ with N for the hydrogen chain in Fig. 14. Localized and optimized orbitals bases give a weight scaling as $\lambda = \mathcal{O}(N^{1.2})$.

APPENDIX D: TROTTER ERROR FOR ELECTRONIC STRUCTURE HAMILTONIANS

In this appendix we report in detail on our numerical analysis of Trotter error. We briefly recall our motivation




FIG. 14. The weight λ of the Hamiltonian depends on its decomposition. Here we show the scaling of λ for the hydrogen chain with system size N_H , for canonical orbitals, localized orbitals, an ansatz based on the eigenvectors of the largest Cholesky vector, and an optimized representation from gradient descent over the space of orbital basis changes.

for doing so. There are many ways to analytically bound the Trotter error. The tightest bounds are based on nested commutators. Nevertheless, even if these have the correct scaling, they are still loose by some constant factor [7]. Additionally, if the Hamiltonian has L terms, the commutator bound for the p th order in principle may require as many as $\mathcal{O}(L^{p+1})$ operations, which is not feasible to compute exactly in practice, even for $p = 2$, for molecular electronic structure Hamiltonians with many orbitals.

For these reasons, we take a numerical approach also taken in Refs. [6,17] and numerically compute the Trotter error for small systems, to then extrapolate to larger systems. Recall that if H is the Hamiltonian with ground state E_0 , $\tilde{H}$ the effective Hamiltonian of the Trotterized time evolution with ground state $\tilde{E}_0$, $U(\delta)$ is the exact time evolution for time δ and $S_p(\delta)$ is a Trotter step with time δ , then we have

$$|E_0 - \tilde{E}_0| \leq C_{\text{gs}} \delta^p \quad \text{and} \quad \|U(\delta) - S_p(\delta)\|_{\text{op}} \leq C_{\text{op}} \delta^{p+1}, \quad (\text{D1})$$

for an order p product formula. Here, we let C_{gs} and C_{op} be the smallest constants for which the bound holds. Crucially, these values depend on the Hamiltonian, as well as the system size. The cost of the product formula scales with $C_{\text{gs}}^{1/p}$.

We estimate C_{gs} and C_{op} by numerically computing $U(\delta)$ and $S_p(\delta)$ for a range of values of δ , and fit a power law to the resulting ground-state energy and operator norm errors. We find that for a wide range of δ we get an excellent fit (even though in principle for large δ there will

be higher order corrections which are dominant) and use this to determine C_{gs} and C_{op} , leading to Fig. 9.*** For the numerical examples, we use a symmetry-preserving Bravyi-Kitaev mapping [86]. The Trotter error depends on the ordering of the terms; we choose a lexicographic ordering. See Refs. [87,88] for a numerical study of the effect of the choice of fermion-to-qubit mapping and ordering on the Trotter error.

In Fig. 15 we compare C_{gs} and C_{op} for the small molecule benchmark set and the hydrogen chain for $p = 2$. We observe both correlate well (especially C_{op}) with λ . Additionally, we also check the correlation with the commutator

$$\alpha = \sum_{l_1, l_2} \|[H_{l_1}, H_{l_2}]\|_{\text{op}} = \sum_{l_1, l_2} |h_{l_1} h_{l_2}| \|[P_{l_1}, P_{l_2}]\| \quad (\text{D2})$$

and find a comparable correlation.

One can also use higher-order product formulas, which leads to a better scaling in t , the time to be simulated. However, at fixed finite time t it is not clear whether this gives an improvement, since increasing the order also increases the number of stages N_{stage} of the product formula, and N_{stage} increases exponentially with the order p for the Trotter-Suzuki product formulas. In Fig. 16 we compare the second and fourth order product formulas for the same benchmark systems. Here we plot the factor that determines the difference in cost, which is $N_{\text{stage}} C_{\text{gs}}^{1/p}(p, \{H_l\}) \varepsilon^{-1/p}$. The number of stages is $N_{\text{stage}} = 2, 10$ for $p = 2, 4$ and we take $\varepsilon = 0.0016$. We do not observe a consistent improvement from using the fourth order method. Of course, for sufficiently small accuracy ε , the higher-order method will eventually be superior. However, for this work we conclude that it suffices to use second order (which has the added benefit that it does need to evolve backward in time, which is beneficial for the partially randomized product formula). In Fig. 16 we compute the Trotter constants for the Pauli and factorized representation of the Hamiltonian for the hydrogen chain, to estimate the scaling with N .

We also briefly comment on the first-order method: for the ground-state energy it has *second-order scaling*. This is an easy consequence of the fact that in perturbation theory, the first-order correction $\langle \psi_0 | [H_{l_1}, H_{l_2}] | \psi \rangle$ vanishes [89]. However, since the second-order method only has two stages, the additional cost is not very high.

1. Trotter error for partial randomization

For partially randomized product formulas, given

$$H = \sum_l h_l P_l, \quad (\text{D3})$$

one groups together a subset of terms into H_R . The Trotterization error for the partially randomized method is the




FIG. 15. The left plot shows the Trotter constants for the operator norm C_{op} and the ground-state energy C_{gs} , for the benchmark set of small molecules and its correlation with α [defined in Eq. (D2)]. The right plot shows the same data, compared to the weights λ .

Trotterization error of dividing the Hamiltonian into terms $h_l P_l$ and keeping H_R as a single term. For the performance of the partially random method it is clearly of interest how the Trotter error depends on the choice of H_R . There are two intuitions of what could happen to the Trotter error. Firstly, if many terms are in H_R , we are discretizing less, and the Trotter error should become smaller. In particular, if all terms are in H_R , there is no Trotter error whatsoever. On the other hand, the Trotter error decreases if one takes very small steps. Grouping together many small terms into one larger term H_R means performing a relatively large step in the direction of H_R , which could *increase* the Trotter error.

In Fig. 18 we show numerically that for λ_R relatively small (the regime which is of most interest for partial randomization), the Trotter error remains essentially unchanged. For intermediate values of λ_R , the Trotter error

in fact increases, but only moderately so. When λ_R tends to λ , the Trotter error decreases significantly.

These findings justify our choice to optimize the partially randomized methods using the Trotter error estimate for the fully deterministic method. Indeed, the regime where λ_R is a large fraction of λ will not lead to significant speedup over the fully randomized method (so it does not help much that the Trotter error is smaller), and when λ_R is reasonably small, the Trotter error of the fully deterministic product formula is a very good approximation.

2. Truncation error

In the partially randomized method the small terms are implemented in a randomized fashion. An alternative way




FIG. 16. This figure shows the effect of the choice of order p of the product formula. We plot $\text{Cost} = PC_{\text{gs}}^{1/p} \varepsilon^{-1/p}$ for $p = 2$ and $p = 4$ Trotter-Suzuki product formulas and $\varepsilon = 0.0016$ for the benchmark set of small molecules and for the hydrogen chain.




FIG. 17. For order $p = 2$ the Trotter ground-state error constant $\sqrt{C_{\text{gs}}}$ [defined in Eq. (17)] for the hydrogen chain with N atoms, both for the Pauli representation and the factorized representation.




FIG. 18. Trotter error as a function of the fraction of the randomized part for different active spaces of alanine and Fe_2S_2 . The dashed line indicates the fully deterministic Trotter error, while the darkness of the data points corresponds to the size of the active space, with darker points corresponding to larger active space size. The error is written in terms of C_{gs} , defined in Eq. (17).

to deal with small terms is to simply neglect them altogether, provided the resulting truncation error is smaller than the desired precision. In Fig. 19 we show the dependence of truncation errors on the fraction of remaining terms for the same set of model systems as used in the Trotter error analysis in Appendix D. For the Pauli decomposition we computed the optimal partition into randomized and deterministic parts as described in Appendix E 3. We then computed the error in the ground-state energy when keeping only the terms in the deterministic part. Note that we do not consider errors due to Trotterization here. Figure 19 shows that the deterministic terms typically make up only a fraction of 0.05 or less of the total number of terms. Simply neglecting the other terms results in an error that is one to two orders of magnitude higher than the target precision.

In Fig. 20 we compute exact truncation errors over the whole range of terms for various active space Hamiltonians of alanine. Here we see that the minimal number of terms one could reach by neglecting small terms (such that the truncation error is not more than 0.0016) is about an order of magnitude higher than the number of deterministic terms in the partially randomized method. However, still surprisingly many terms can be truncated (if the exact truncation error is known), here to a fraction of about 0.2. Note that in practice some of the error budget also needs to be assigned to the Trotterization and phase estimation error.

Consequently, the partially randomized method achieves a reduction of the deterministically implemented terms that would not be possible by a simple truncation approach. In general, a downside of trying to (deterministically) truncate as many terms as possible from the Hamiltonian, is




FIG. 19. For the Hamiltonians from the benchmark set, we plot the error in the ground-state energy as a function of the fraction of remaining terms after truncating the ones with smallest weight. We use the truncation which is optimal for the cost of the partially random product formula, so we compare the energies of H and H_D . The red dashed line indicates chemical accuracy, 0.0016.

that it is hard to know a priori to what point one can truncate without incurring significant error, compromising the accuracy guarantees of the energy approximation procedure.

Yet another alternative to deal with small terms in the Hamiltonian is to treat them perturbatively. Quantum algorithms for second-order perturbation theory corrections have been proposed previously [79,90]. However, such an approach requires many samples, and for each sample the ground state of the truncated Hamiltonian needs




FIG. 20. Error in the ground-state energy as a function of the fraction of remaining terms after truncating the ones with smallest weight. The five blue lines correspond to the active space Hamiltonians of alanine with 4, 5, 6, 7, and 8 orbitals (darker lines for larger active spaces). The red dashed line indicates chemical accuracy, 0.0016.

to be prepared, and the inverse of the truncated Hamiltonian needs to be implemented. This potentially negates any cost-advantage from truncating the Hamiltonian, which was the outcome of Ref. [79]. A full analysis of a perturbation theory approach in our context is beyond the scope of this work.

APPENDIX E: DETAILED COST ESTIMATES AND COMPILATION

In this part of the appendix we are giving a detailed account of how the gate counts are estimated for ground-state energy estimation using the deterministic, randomized, or partially randomized simulation methods. For each method we are estimating the number of gate G , where we can either measure the number of (non-Clifford) Toffoli gates, or the number of two-qubit gates. We are given a Hamiltonian H , and wish to estimate the ground-state energy to target precision ε (root-mean-square error). In these estimates we assume access to an exact eigenstate of H , but as discussed this assumption can easily be relaxed. The total cost depends on various parameters, some of which depend on H and some of which can be optimized. We give an overview of the relevant parameters in Table I.

1. Gate counts for Trotterization

The Hamiltonian has a decomposition $H = \sum_{l=1}^L H_l$. We use a product formula of order p . For most of the resource estimates we use second order ($p = 2$). As explained in Sec. III, we perform phase estimation on the Trotter unitary $S_p(\delta) = e^{-i\delta\tilde{H}} \approx e^{-i\delta H}$ for this task, computing the ground energy $\tilde{E}_0$ of the effective Hamiltonian $\tilde{H}$. The error between the ground states of H and $\tilde{H}$ behaves as $\varepsilon_{\text{trot}} = |E_0 - \tilde{E}_0| \leq C_{\text{gs}}\delta^p$. Additionally, when applying phase estimation, we obtain $\tilde{E}_0$ with mean square error ε_{qpe} . This error is unbiased and independent of the Trotter error, which means that in order get an estimate of precision ε for E_0 we should have

$$\varepsilon^2 = \varepsilon_{\text{qpe}}^2 + \varepsilon_{\text{trot}}^2. \quad (\text{E1})$$

To reach precision ε_{qpe} for $\tilde{E}_0$, we need to estimate the phase of $S_p(\delta)$ to precision $\varepsilon_{\text{qpe}}\delta$. We estimated in Appendix B that robust phase estimation needs on average $5\pi/(\varepsilon_{\text{qpe}}\delta)$ implementations of $S_p(\delta)$ to do so. The number of rounds in the phase estimation is

$$M = \left\lceil \log_2 \left(\frac{0.08\pi}{\varepsilon_{\text{qpe}}\delta} \right) \right\rceil. \quad (\text{E2})$$

Let G_{det} be the average gate cost to implement a single evolution along the H_l , then we see that the total gate cost

is given by

$$G = 5\pi N_{\text{stage}} L G_{\text{det}} \frac{C_{\text{gs}}^{1/p}}{\varepsilon_{\text{qpe}} \varepsilon_{\text{trot}}^{1/p}}. \quad (\text{E3})$$

Minimizing the cost under the error budget in Eq. (E1) results in $\varepsilon_{\text{qpe}} = \varepsilon \sqrt{p/(p+1)}$, and trotter step size and optimal total cost

$$\delta = \left(\frac{\varepsilon}{C_{\text{gs}} \sqrt{p+1}} \right)^{1/p}, \quad (\text{E4})$$

$$G = 5\pi N_{\text{stage}} L G_{\text{det}} \sqrt{\frac{(p+1)^{1+1/p}}{p}} \frac{C_{\text{gs}}^{1/p}}{\varepsilon^{1+1/p}}. \quad (\text{E5})$$

We can reduce the cost by factor 2 using the method in Appendix A 4. At this point, the cost is agnostic of the gate set used. For the Toffoli count we additionally need to approximate with a gate set of Cliffords and Toffolis and consider the resulting synthesis error. Here we follow the analysis of Refs. [17,47]. We decompose circuits into Cliffords plus arbitrary single-qubit rotations, which can be synthesized to error ε' at a cost of

$$G_{\text{rot}} = 1.14 \log_2 \left(\frac{1}{\varepsilon'} \right) + 9.2. \quad (\text{E6})$$

T gates [91]. We convert from Toffoli to T gates using 1 Toffoli gate for every 2 T gates [92]. Let n_{rot} be the number of single-qubit rotations per trotter stage, each approximated to error ε' . We then want to bound the synthesis error $\varepsilon_{\text{synth}}$, i.e., the difference in the ground state energies of the exact effective Hamiltonian and the effective Hamiltonian with imperfectly synthesized single-qubit rotations, denoted by H' . One can bound [17]

$$\|e^{-i\delta\tilde{H}} - e^{-i\delta H'}\| \leq N_{\text{stage}} n_{\text{rot}} \varepsilon', \quad (\text{E7})$$

so to first order in δ the ground-state error is bounded by

$$|\tilde{E}_0 - E'_0| \leq \|\tilde{H} - H'\| \leq N_{\text{stage}} n_{\text{rot}} \varepsilon' \delta^{-1}. \quad (\text{E8})$$

Consequently, to keep the synthesis error below $\varepsilon_{\text{synth}}$, we choose

$$\varepsilon' = \frac{\delta \varepsilon_{\text{synth}}}{N_{\text{stage}} n_{\text{rot}}} = \frac{\varepsilon_{\text{synth}}}{N_{\text{stage}} n_{\text{rot}}} \left(\frac{\varepsilon}{C_{\text{gs}} \sqrt{p+1}} \right)^{1/p}. \quad (\text{E9})$$

We do not optimize the total error budget with respect to $\varepsilon_{\text{synth}}$, instead, when computing the Toffoli count, we allocate $\varepsilon = 1.5$ mH and $\varepsilon_{\text{synth}} = 0.1$ mH. Alternatively, if we measure two-qubit gates, we ignore the synthesis error, we let $\varepsilon = 1.6$ mH and set G_{det} to be the number of 2-qubit gates.

So far, we have assumed a general Hamiltonian $H = \sum_l H_l$. We now specialize to specific representations.

TABLE 1. Relevant parameters for gate counts for of phase estimation of ground-state energies using deterministic and (partially) randomized product formulas.

Parameter	Notation
Phase estimation parameters	
ε	Target precision for the energy.
M	Number of rounds in phase estimation.
N_m	Number of samples in round m , $N_m = N_M + D(M - m)$.
Deterministic product formulas	
L	Number of terms in the Hamiltonian.
$\varepsilon_{\text{qpe}}, \varepsilon_{\text{trot}}$	Dividing error budget between phase estimation and Trotter error.
δ	Trotter step size.
p	Order product formula.
N_{stage}	Number of stages product formula.
C_{gs}	Trotter error constant for the ground-state energy.
G_{det}	Number of elementary gates to perform one elementary evolution $\exp(-i\delta H_l)$.
Randomized product formulas	
λ	Weight of the Hamiltonian.
τ	Step size
φ	Rotation angle $\varphi = \arctan(\tau)$.
G_{rand}	Number of elementary gates to time evolve a single Pauli operator P_l .
B	Damping factor from the randomized product formula.
Partially randomized product formulas	
L_D	Number of terms treated deterministically.
λ_R	Weight of the terms treated randomly.
κ	Scaling factor for the number of rotations for randomized product formula.

a. Pauli Hamiltonian

We start with the case where the Hamiltonian has a representation as a linear combination of Pauli operators, where $H_l = h_l P_l$, as in Eq. (C7). The time evolution of a single term is then a Pauli rotation, $V_l(\delta h_l) = \exp(-i\delta h_l P_l)$. This can be implemented using a single-qubit rotation and $2(|P_l| - 1)$ CNOT gates to implement. Here we use $|P_l|$ to denote the number of qubits on which P_l acts nontrivially.

The number of Pauli rotations per stage simply equals the number of terms L . The worst-case scaling of the number of terms with the number of orbitals is $L = \mathcal{O}(N^4)$, but in practice many terms are sufficiently small to be truncated (as discussed in Appendix C) and H is much more sparse. This depends on the choice of single orbital basis as well. We find that minimizing the weight λ through orbital rotations and symmetry shifts typically also leads to more sparse representations of H . For our estimates, we truncate the smallest coefficients of H , such that the total weight $\sum_{l>L'} |h_l|$ of truncated terms is at most a small threshold 10^{-3} .

When counting two-qubit gates we can contract the innermost CNOT gates together with the single-qubit rotation into a single 2-qubit gate, so the number of two-qubit gates per Pauli rotation is $2|P_l| - 3$. Therefore the number of 2-qubit gates per operation is $G_{\text{det}} = 2|\bar{P}| - 3$, where $|\bar{P}|$ denotes the average Pauli weight. Reductions in two-qubit gate cost are achieved by reducing L or by

reducing the average Pauli weight. Again, minimizing λ also reduces L . The average Pauli weight depends on the choice of fermion-to-qubit mapping. On N spin-orbitals the Bravyi-Kitaev mapping gives an average Pauli weight of $\mathcal{O}(\log(N))$, compared to $\mathcal{O}(N)$ for the Jordan-Wigner and parity mappings. We use a symmetry-conserving Bravyi-Kitaev mapping, which has the additional advantage of reducing the number of qubits by 2 by taking into account spin- and particle-number conservation [86]. Yet another degree of freedom is the choice of term ordering the ordering in a Trotter stage, which can lead to significant cancellations of 2-qubit gates [93]. We use the lexicographic ordering, which has a relatively large number of cancellations [94].

b. Factorized Hamiltonian

In this case, we have a Hamiltonian of the form

$$H = T + \sum_{l=1}^L (U^{(l)})^\dagger V^{(l)} U^{(l)}, \quad V^{(l)} = \sum_{\sigma} \sum_{p=1}^{\rho_l} \lambda_p^{(l)} n_{p\sigma}. \quad (\text{E10})$$

We first discuss the two-body terms. In order to implement one stage in a Trotter step, we need to implement L basis changes $U^{(l)} (U^{(l-1)})^\dagger$ (which can be combined into a single orbital basis change) and evolutions of the form $\exp(-i\delta V^{(l)})$. This can be done [24,64] using fermionic swap networks. The orbital basis rotation $U^{(l)}$ for both

spin values can be implemented with $2 \left(\binom{N}{2} - \binom{N-\rho_l}{2} \right) = \mathcal{O}(\rho_l N)$. Givens rotations (which are two-qubit gates and can be implemented using two arbitrary single-qubit gates) and $\exp(-i\delta V^{(l)})$ can be implemented using $\binom{2\rho_l}{2}$ two-qubit gates (requiring one arbitrary single-qubit gate). This gives a total number of single-qubit rotations

$$\begin{aligned} & \sum_{l=1}^L 4 \left(\binom{N}{2} - \binom{N-\rho_l}{2} \right) + \binom{2\rho_l}{2} \\ &= \sum_{l=1}^N 4N(\rho_l - 1) + \rho_l. \end{aligned} \quad (\text{E11})$$

The evolution along T is a single orbital basis change [requiring $2\binom{N}{2}$ Givens rotations]. For this approach to be useful, one has to truncate small $\lambda_p^{(l)}$. For Fig. 10, where we compute the cost of this approach for the hydrogen chain, we truncate such that for each l , if we truncate $\lambda_p^{(l)}$ for $p > p_l$, we have $\sum_{p > p_l} |\lambda_p^{(l)}| \leq \varepsilon'$, where we choose $\varepsilon' = 10^{-3}$.

2. Gate counts for randomized product formulas

For randomized product formulas, the Hamiltonian is assumed to be given in a Pauli representation. The cost is determined by the weight λ and the target precision ε . We are given a Pauli Hamiltonian with some value of λ , and target precision ε . As discussed in Appendix B, we have to perform time simulation for $t = 2^m$ and take N_m samples from a Hadamard test, with $m = 1, \dots, M$ for $M = \lceil \log_2(\lambda \varepsilon^{-1}) \rceil$, which gives a number of Pauli rotations as in Eq. (B17) as

$$r = \frac{16.3\lambda^2}{\varepsilon^2}. \quad (\text{E12})$$

Note that for qDRIFT the number of rotations may be reduced by a factor of 2 (see the discussion at the end of Appendix B 1). Counting two-qubit gates is similar to the previous section. When counting Toffoli gates it turns out that it is helpful that all (or most) rotations are over the same angle, leading to a significantly smaller overhead in converting the number of rotations to the number of Toffoli gates.

The below approach combines ideas from Refs. [13, 21, 22, 47] in the analysis of the Toffoli gate requirements. To begin with, for qDRIFT, in every circuit *all* rotations are over the same angle. For RTE, the angle φ_n depends on the order n sampled of the term in the Taylor series. For small τ , the probability of sampling a term which is not of the lowest order scales with $\mathcal{O}(\tau^2)$. When performing time simulation for time t with $\tau = \Omega(t^{-1})$ and number of steps $r = \mathcal{O}(t^2)$ this means that we only expect $\mathcal{O}(1)$ Pauli rotations with an angle different than φ_0 . This means that

we expect long sequences of equiangular Pauli rotations which dominate the cost.

Additionally, any individual Pauli operator P_l with nonzero h_l in the Hamiltonian commutes with *most* of the other terms; it only anticommutes with $\mathcal{O}(N^3)$ out of $\mathcal{O}(N^4)$ terms. For this reason, one expects relatively long sequences of commuting Pauli rotations in both qDRIFT and RTE, as observed in Ref. [13] [if the terms were selected uniformly, one would have average length $\mathcal{O}(\sqrt{N})$]. Given such a sequence, $e^{i\varphi P_{l_1}} \dots e^{i\varphi P_{l_K}}$ of length $K < 2N$, we may apply a Clifford to map

$$\{P_{l_1}, \dots, P_{l_K}\} \rightarrow \{Z_1, \dots, Z_K\} \quad (\text{E13})$$

to single-qubit Z -rotations on the first K qubits. We may now use Hamming weight phasing [21, 47, 95] to implement these rotations, as suggested for RTE in Ref. [13].

To reduce the cost even further, we can use a single catalyst phase state to implement the rotations [4, 22, 33]. We may choose the step size in the randomized product formula to our liking; in particular we may choose such that the rotations are of angle $2^{-J}\pi$ for integer J . Since we also want to choose the angle such that $\tau = \mathcal{O}(\varepsilon \lambda^{-1})$, we take $J = \lceil \log_2(\lambda \pi^{-1} \varepsilon^{-1}) \rceil$. Let

$$|\phi\rangle = \frac{1}{2^{J/2}} (|0\rangle + e^{2\pi i 2^{-J}} |1\rangle)^{\otimes J} = \frac{1}{2^{J/2}} \sum_{j=0}^{2^J-1} e^{2\pi i j 2^{-J}} |j\rangle \quad (\text{E14})$$

be a phase state on J qubits. At the start of the computation, we initialize the ancilla phase estimation qubit, together with J qubits, in the state

$$\begin{aligned} & \frac{1}{\sqrt{2}} (|0\rangle |+\rangle^{\otimes J} + |1\rangle |\phi\rangle) = (I \otimes \text{QFT}) \frac{1}{\sqrt{2}} (|0\rangle |0\rangle \\ & + |1\rangle |1\rangle), \end{aligned} \quad (\text{E15})$$

where we label the basis for the J qubits by $\{0, 1, \dots, 2^J - 1\}$. It (once) requires T gates or Toffoli gates to prepare this state (a constant cost which is negligible compared to the other contributions). The phase state $|\phi\rangle$ is such that if we add another register with $W \leq J$, and implement addition $|w\rangle |j\rangle \mapsto |w\rangle |j + w\rangle$, then

$$|w\rangle |\phi\rangle \mapsto e^{-2\pi i w 2^{-J}} |w\rangle |\phi\rangle. \quad (\text{E16})$$

We can use J Toffoli gates, together with J ancilla qubits to perform the addition. The state $|+\rangle^{\otimes J}$ is invariant under addition, so we get that

$$\begin{aligned} & \alpha |0\rangle |w\rangle |+\rangle^{\otimes J} + \beta |1\rangle |w\rangle |\phi\rangle \mapsto \alpha |0\rangle |w\rangle |+\rangle^{\otimes J} \\ & + \beta e^{-2\pi i w 2^{-J}} |1\rangle |w\rangle |\phi\rangle. \end{aligned} \quad (\text{E17})$$

To implement $\prod_{k=1}^K e^{i\pi 2^{-J} Z_k}$, on the first K qubits we first compute the Hamming weight on a register of

$W = \lceil \log_2(K+1) \rceil$ qubits. We then apply the above trick to apply (controlled on the QPE ancilla) $e^{i w \pi 2^{-J}}$ to Hamming weight w , and we uncompute the Hamming weight. Note that if $|k\rangle$ has Hamming weight w , then

$$\left(\prod_{k=1}^K e^{i \pi 2^{-J} Z_k} \right) |k\rangle = e^{-i \pi (2w-K) 2^{-J}} |k\rangle, \quad (\text{E18})$$

so we have implemented the correct transformation up to a global phase $e^{-i \pi K 2^{-J}}$ which we can account for in classical postprocessing of the Hadamard test outcomes.

Computing and uncomputing the Hamming weight requires at most $K-1$ Toffoli gates applying the phase to the Hamming weight requires $J-1$ Toffoli gates. This means we use $1 + (J-2)/K$ Toffoli gates per controlled Pauli rotation. The adders require respectively $(K-1)$ and $(J-1)$ ancilla qubits, giving in total $K+2J-2$ ancilla qubits.

For example, for FeMoco, we may take $K=10$ (choosing K not too large, in order to make sure that with high probability we get sequences of commuting terms of length K) and $J=17$, giving an average of 2.5 Toffoli gates per controlled Pauli rotation, and using 42 ancilla qubits.

3. Gate counts for partially randomized product formulas

We assume a given decomposition of H in L_D deterministic terms, and H_R with weight λ_R ,

$$H = \sum_{l=1}^{L_D} H_l + H_R. \quad (\text{E19})$$

The Trotter error concerns the decomposition into H as a sum of the H_l and treating H_R as a single term. For our estimates we use the second-order Trotter-Suzuki product formula (so $p=2$ with $N_{\text{stage}}=2$). It is straightforward to extend to higher orders. Let G_{det} and G_{rand} be the (average) number of Toffolis used per unitary $\exp(-i\delta H_l)$ in the deterministic part, and per Pauli rotation for the randomized part of the Hamiltonian simulation. These can be different.

We choose the parameters ε_{qpe} and $\varepsilon_{\text{trot}}$, the time step δ for the deterministic product formula, and the number of rounds M for phase estimation as in Eqs. (E1), (E2), and (E4). For the compilation of the randomized part, it is convenient to choose time step τ such that we can use Pauli rotations over angles $\pi 2^{-J}$. For $m=1, \dots, M$ we use $2N_m$ times the Hadamard tests with the partially randomized product formula to estimate (see Lemma A3),

$$B^{-1} \langle \psi | S_2(\delta)^{2^m} | \psi \rangle. \quad (\text{E20})$$

As explained at the end of Sec. B1, for RTE alone it is optimal to choose a total of $r = 2t^2 = 2\lambda_R^2 \delta^2 2^{2m}$ Pauli

rotations. For the partially randomized method it is advantageous to allow for r to vary, so we introduce a factor $\kappa \geq 0$ and set $r = \kappa \lambda_R^2 \delta^2 2^{2m}$, which gives a damping factor of $B \leq e^{1/\kappa}$. The number of samples N_m scales with the damping factor B^2 , as discussed in Appendix B. We use the observation in Appendix A4 to reduce the evolution time of the deterministic part by a factor of 2, so the deterministic product formula needs 2^{m-1} rather than 2^m applications of $S_2(\delta)$. This gives a total number of gates,

$$G = \sum_{m=0}^M 2N_m (G_{\text{det}} N_{\text{stage}} L_D 2^{m-1} + G_{\text{rand}} \kappa \lambda_R^2 \delta^2 2^{2m}). \quad (\text{E21})$$

Note that if we take $L_D=0$, this reduces (up to possibly different rounding) to the cost of phase estimation to precision ε_{qpe} using a randomized product formula, independent of the choice of δ . If we choose $N_m = N_M + D(M-m) = B^2(11+4(M-m)) = e^{2/\kappa}(11+4(M-m))$, the total number gates is bounded as in Eqs. (B30) and (B31) by

$$\begin{aligned} G &\leq 2G_{\text{det}} N_{\text{stage}} L_D (N_M + D) 2^M + 8G_{\text{rand}} \left(\frac{N_M}{3} + \frac{D}{9} \right) \\ &\quad \kappa \lambda_R^2 \delta^2 2^{2M} \\ &= 30G_{\text{det}} N_{\text{stage}} L_D e^{2/\kappa} \frac{0.1\pi}{\delta \varepsilon_{\text{qpe}}} + \frac{296}{9} G_{\text{rand}} \kappa e^{2/\kappa} \frac{(0.1\pi \lambda_R)^2}{\varepsilon_{\text{qpe}}^2}. \end{aligned} \quad (\text{E22})$$

The optimal choice of κ solves a quadratic equation. Again, we require the total error to be bounded by

$$\varepsilon^2 \leq \varepsilon_{\text{qpe}}^2 + \varepsilon_{\text{trot}}^2 = \varepsilon_{\text{qpe}}^2 + C_{\text{gs}}^2 \delta^4 \quad (\text{E23})$$

such that we substitute $\varepsilon_{\text{qpe}} = \sqrt{\varepsilon^2 - C_{\text{gs}}^2 \delta^4}$ in Eq. (E22). Finally, we find the minimal cost by optimizing for L_D (which determines λ_R). In the following we discuss two strategies for decomposing the Hamiltonian into deterministic and randomized parts, and the associated gate counts.

a. Pauli decomposition

A natural scheme for the partially randomized method is to start from the Hamiltonian in the Pauli sum decomposition after the fermion-to-qubit mapping, $H = \sum_{l=1}^L h_l P_l$. We decompose into H_D and H_R by assigning the L_D terms with largest weight h_l to H_D , which makes it straightforward to find the optimal choice of L_D based on Eq. (E22).

If we count two-qubit gates, as discussed for deterministic product formulas, the terms in the deterministic part can be ordered lexicographically in order to achieve cancellations of CNOT gates.




FIG. 21. Toffoli cost for partially randomized product formulas with respect to varying N and ε , with a fit to a scaling $\mathcal{O}(N^2 \varepsilon^{-1.7})$.

If we count Toffoli gates we can again apply the modified Hamming weight phasing to the randomized part, as presented in Appendix E 2. We can additionally leverage this to reduce the gate synthesis cost of the deterministic part. The idea is that while in principle the deterministic part consists of terms $h_l P_l$ with arbitrary real coefficients h_l , we may round h_l to a desired value and assign the remainder to H_R . This slightly increases the value of λ_R , but may significantly reduce the synthesis cost of time evolving along P_l in the deterministic part. This approach is similar to previous methods using randomization for compiling arbitrary-angle rotations [57–59]. We obtain further reductions by expanding the rounded h_l bitwise and reorder terms so we can apply Hamming weight phasing.

We now explain this idea in detail. Consider the binary expansion of the rotation angle of the l th term

$$\frac{h_l \delta}{\pi} = \sum_{J \geq 1} x_{lJ} 2^{-J} \quad \text{with } x_{lJ} \in \{0, 1\}. \quad (\text{E24})$$

Grouping bits of the same significance and truncating to n_{max} bits of precision gives

$$\delta H = \sum_{J=1}^{n_{\text{bit}}} \underbrace{\sum_{l=1}^{L_J} \pi 2^{-J} P_l}_{\delta H_J} + \delta H_{\text{rem}}, \quad (\text{E25})$$

where H_{rem} contains the remaining higher bits of precision. For the deterministic product formula, we now apply all the terms in H_J consecutively. This consists of L_J Pauli rotations with the same angle, so we can use Hamming weight phasing. Moreover, for the deterministic part we are free to choose the order of the Pauli rotations, and since any individual term anticommutes with $\mathcal{O}(N^3)$ out of the total $\mathcal{O}(N^4)$ terms, one can group into long sequences of commuting terms. While decomposing into

the H_J increases the number of terms, they are cheaper to implement, reducing the total cost.

We now comment on two subtleties which influence this approach. Firstly, one needs to group the terms in each H_J into commuting groups. Such groupings have been studied extensively for the sake of reducing measurement overheads when measuring energies [96,97]. While one can typically find quite large groupings, this can be computationally intensive. For our compilation purposes, if we have groups of K commuting Pauli operators, implementing a rotation over angle $\pi 2^{-J}$ requires $1 + (J - 2)/K$ Toffoli gates per Pauli. This means that groups of moderate size suffice to have a low overhead, which can be constructed using simple heuristics. As a lower bound, which we use for our numerical estimates, when considering Eq. (C6), it is not hard to see that it is possible to group into groups of 12 commuting Pauli operators. The second subtlety is that the procedure of decomposing each term into several bits and reordering the terms will in principle affect the Trotter error. The effect of ordering on Trotter error has been studied numerically for molecular electronic structure Hamiltonians [88]. We expect that decomposing by bits of precision will not increase the Trotter error (since it decomposes into shorter time evolutions, so this should be beneficial). For simplicity we assume that the Trotter error is the same as for the fully deterministic product formula and leave a detailed numerical exploration to future work. For FeMoco, we find a total cost of 1.5×10^{12} Toffoli gates using this strategy, as shown in Fig. 1. Note that this approach increases the number of ancilla qubits, since we need to prepare phase states for different bits of precision J .

b. Scaling for the hydrogen chain

As a benchmark system we use the hydrogen chain. The result for $\varepsilon = 0.0016$ is shown in Fig. 2. We can also vary over ε as well, we can fit an expression of the form $C N^a \varepsilon^{-b}$




FIG. 22. Two-qubit gate count in terms of CNOT required for the hydrogen chain for precision $\varepsilon = 0.0016$.

to the resulting cost. In Fig. 21 we show scalings with respect to N and ε (keeping the other variable fixed). We find a fit for the Toffoli cost of $\mathcal{O}(N^2\varepsilon^{-1.7})$.

In Fig. 22 we show the number of CNOT gates required for phase estimation using deterministic, randomized and partially randomized product formulas. Note that this has a different scaling than the Toffoli count, since the number of two-qubit gates required per Pauli rotation scales with N . As for the Toffoli gate count, we observe a significant improvement using partial randomization.

c. Double-factorized decomposition

If we truncate to L_D terms in the factorization, where the l th term has rank at most ρ_l , then for the deterministic part we have cost as described in Appendix E 1 b. The cost of the randomized part remains the same. The performance of this approach for the hydrogen chain is shown in Fig. 10.

[1] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su, Toward the first quantum simulation with quantum speedup, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9456 (2018).
 [2] Vera von Burg, Guang Hao Low, Thomas Häner, Damian S. Steiger, Markus Reiher, Martin Roetteler, and Matthias Troyer, Quantum computing enhanced computational catalysis, *Phys. Rev. Res.* **3**, 033055 (2021).
 [3] Yuan Su, Dominic W. Berry, Nathan Wiebe, Nicholas Rubin, and Ryan Babbush, Fault-tolerant quantum simulations of chemistry in first quantization, *PRX Quantum* **2**, 040332 (2021).
 [4] Joonho Lee, Dominic W. Berry, Craig Gidney, William J. Huggins, Jarrod R. McClean, Nathan Wiebe, and Ryan Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *PRX Quantum* **2**, 030305 (2021).

[5] Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan, Quantum algorithms for quantum chemistry and quantum materials science, *Chem. Rev.* **120**, 12685 (2020).
 [6] Ryan Babbush, Jarrod McClean, Dave Wecker, Alán Aspuru-Guzik, and Nathan Wiebe, Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation, *Phys. Rev. A* **91**, 022311 (2015).
 [7] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu, Theory of Trotter error with commutator scaling, *Phys. Rev. X* **11**, 011020 (2021).
 [8] Guang Hao Low, Yuan Su, Yu Tong, and Minh C. Tran, Complexity of implementing Trotter steps, *PRX Quantum* **4**, 020323 (2023).
 [9] Yuan Su, Hsin-Yuan Huang, and Earl T. Campbell, Nearly tight trotterization of interacting electrons, *Quantum* **5**, 495 (2021).
 [10] Guang Hao Low and Isaac L. Chuang, Optimal Hamiltonian simulation by quantum signal processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
 [11] Guang Hao Low and Isaac L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
 [12] Earl Campbell, Random compiler for fast Hamiltonian simulation, *Phys. Rev. Lett.* **123**, 070503 (2019).
 [13] Kianna Wan, Mario Berta, and Earl T. Campbell, Randomized quantum algorithm for statistical phase estimation, *Phys. Rev. Lett.* **129**, 030503 (2022).
 [14] Yingkai Ouyang, David R. White, and Earl T. Campbell, Compilation by stochastic Hamiltonian sparsification, *Quantum* **4**, 235 (2020).
 [15] Shi Jin and Xiantao Li, A partially random Trotter algorithm for quantum Hamiltonian simulations, *Commun. Appl. Math. Comput.* **7**, 442 (2025).
 [16] Matthew Hagan and Nathan Wiebe, Composite quantum simulations, *Quantum* **7**, 1181 (2023).
 [17] Markus Reiher, Nathan Wiebe, Krysta M. Svore, Dave Wecker, and Matthias Troyer, Elucidating reaction mechanisms on quantum computers, *Proc. Natl. Acad. Sci. U.S.A.* **114**, 7555 (2017).
 [18] Dominic W. Berry, Craig Gidney, Mario Motta, Jarrod R. McClean, and Ryan Babbush, Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization, *Quantum* **3**, 208 (2019).
 [19] Dario Rocca, Cristian L. Cortes, Jerome Gonthier, Pauline J. Ollitrault, Robert M. Parrish, Gian-Luca Anselmetti, Matthias Degroote, Nikolaj Moll, Raffaele Santagati, and Michael Streif, Reducing the runtime of fault-tolerant quantum simulations in chemistry through symmetry-compressed double factorization, *J. Chem. Theory Comput.* **20**, 4639 (2024).
 [20] Guang Hao Low, Robbie King, Dominic W. Berry, Qiushi Han, A. Eugene DePrince III, Alec White, Ryan Babbush, Rolando D. Somma, and Nicholas C. Rubin, Fast quantum simulation of electronic structure by spectrum amplification *Phys. Rev. X* **15**, 041016 (2025).
 [21] Craig Gidney, Halving the cost of quantum addition, *Quantum* **2**, 74 (2018).
 [22] Yuval R. Sanders, Dominic W. Berry, Pedro C. S. Costa, Louis W. Tessler, Nathan Wiebe, Craig Gidney, Hartmut Neven, and Ryan Babbush, Compilation of fault-tolerant quantum heuristics for combinatorial optimization, *PRX Quantum* **1**, 020312 (2020).

- [23] Dominic W. Berry, Yu Tong, Tanuj Khattar, Alec White, Tae In Kim, Sergio Boixo, Lin Lin, Seunghoon Lee, Garnet Kin Chan, Ryan Babbush, *et al.*, Rapid initial state preparation for the quantum simulation of strongly correlated molecules [PRX Quantum 6, 020327 \(2025\)](#).
- [24] Mario Motta, Erika Ye, Jarrod R. McClean, Zhendong Li, Austin J. Minnich, Ryan Babbush, and Garnet Kin-Lic Chan, Low rank representations for quantum simulation of electronic structure, [npj Quantum Inf. 7, 83 \(2021\)](#).
- [25] Ian D. Kivlichan, Christopher E. Granade, and Nathan Wiebe, Phase estimation with randomized Hamiltonians, [arXiv:1907.10070](#).
- [26] Abhishek Rajput, Alessandro Roggero, and Nathan Wiebe, Hybridized methods for quantum simulation in the interaction picture, [Quantum 6, 780 \(2022\)](#).
- [27] Oriel Kiss, Michele Grossi, and Alessandro Roggero, Importance sampling for stochastic quantum simulations, [Quantum 7, 977 \(2023\)](#).
- [28] Norm M. Tubman, Carlos Mejuto-Zaera, Jeffrey M. Epstein, Diptarka Hait, Daniel S. Levine, William Huggins, Zhang Jiang, Jarrod R. McClean, Ryan Babbush, Martin Head-Gordon, *et al.*, Postponing the orthogonality catastrophe: Efficient state preparation for electronic structure simulations on quantum devices, [arXiv:1809.05523](#).
- [29] Seunghoon Lee, Joonho Lee, Huanchen Zhai, Yu Tong, Alexander M. Dalzell, Ashutosh Kumar, Phillip Helms, Johnnie Gray, Zhi-Hao Cui, Wenyuan Liu, *et al.*, Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry, [Nat. Commun. 14, 1952 \(2023\)](#).
- [30] Stepan Fomichev, Kasra Hejazi, Modjtaba Shokrian Zini, Matthew Kiser, Joana Fraxanet Morales, Pablo Antonio Moreno Casares, Alain Delgado, Joonsuk Huh, Arne-Christian Voigt, Jonathan E. Mueller, *et al.*, Initial state preparation for quantum chemistry on quantum computers, [PRX Quantum 5, 040339 \(2024\)](#).
- [31] Mihael Erakovic, Freek Witteveen, Dylan Harley, Jakob Günther, Moritz Bensberg, Oinam Romesh Meitei, Minsik Cho, Troy Van Voorhis, Markus Reiher, and Matthias Christandl, High ground state overlap via quantum embedding methods, [PRX Life 3, 013003 \(2025\)](#).
- [32] Maximilian Mörchen, Guang Hao Low, Thomas Weymuth, Hongbin Liu, Matthias Troyer, and Markus Reiher, Classification of electronic structures and state preparation for quantum computation of reaction chemistry, [arXiv:2409.08910](#).
- [33] A. Y. Kitaev, Quantum measurements and the Abelian stabilizer problem, [arXiv:quant-ph/9511026](#).
- [34] Miroslav Dobšíček, Göran Johansson, Vitaly Shumeiko, and Göran Wendin, Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark, [Phys. Rev. A 76, 030306\(R\) \(2007\)](#).
- [35] Shelby Kimmel, Guang Hao Low, and Theodore J. Yoder, Robust calibration of a universal single-qubit gate set via robust phase estimation, [Phys. Rev. A 92, 062315 \(2015\)](#).
- [36] Nathan Wiebe and Chris Granade, Efficient Bayesian phase estimation, [Phys. Rev. Lett. 117, 010503 \(2016\)](#).
- [37] Rolando D. Somma, Quantum eigenvalue estimation via time series analysis, [New J. Phys. 21, 123025 \(2019\)](#).
- [38] Lin Lin and Yu Tong, Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers, [PRX Quantum 3, 010318 \(2022\)](#).
- [39] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone, Quantum metrology, [Phys. Rev. Lett. 96, 010401 \(2006\)](#).
- [40] Zhiyan Ding and Lin Lin, Even shorter quantum circuit for phase estimation on early fault-tolerant quantum computers with applications to ground-state energy estimation, [PRX Quantum 4, 020331 \(2023\)](#).
- [41] B. L. Higgins, D. W. Berry, S. D. Bartlett, M. W. Mitchell, H. M. Wiseman, and G. J. Pryde, Demonstrating Heisenberg-limited unambiguous phase estimation without adaptive measurements, [New J. Phys. 11, 073023 \(2009\)](#).
- [42] Federico Belliardo and Vittorio Giovannetti, Achieving Heisenberg scaling with maximally entangled states: An analytic upper bound for the attainable root-mean-square error, [Phys. Rev. A 102, 042613 \(2020\)](#).
- [43] Hongkang Ni, Haoya Li, and Lexing Ying, On low-depth algorithms for quantum phase estimation, [Quantum 7, 1165 \(2023\)](#).
- [44] Thomas E. O'Brien, Brian Tarasinski, and Barbara M. Terhal, Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments, [New J. Phys. 21, 023022 \(2019\)](#).
- [45] David Poulin, Matthew B. Hastings, Dave Wecker, Nathan Wiebe, Andrew C. Doherty, and Matthias Troyer, The Trotter step size required for accurate quantum simulation of quantum chemistry, [QIC 15, 361 \(2015\)](#).
- [46] Dave Wecker, Bela Bauer, Bryan K. Clark, Matthew B. Hastings, and Matthias Troyer, Gate-count estimates for performing quantum chemistry on small quantum computers, [Phys. Rev. A 90, 022305 \(2014\)](#).
- [47] Ian D. Kivlichan, Craig Gidney, Dominic W. Berry, Nathan Wiebe, Jarrod McClean, Wei Sun, Zhang Jiang, Nicholas Rubin, Austin Fowler, Alán Aspuru-Guzik, *et al.*, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization, [Quantum 4, 296 \(2020\)](#).
- [48] Burak Şahinoğlu and Rolando D. Somma, Hamiltonian simulation in the low-energy subspace, [npj Quantum Inf. 7, 119 \(2021\)](#).
- [49] Chi-Fang Chen and Fernando G. S. L. Brandão, Average-case speedup for product formulas, [Commun. Math. Phys. 405, 32 \(2024\)](#).
- [50] Qi Zhao, You Zhou, Alexander F. Shaw, Tongyang Li, and Andrew M. Childs, Hamiltonian simulation with random inputs, [Phys. Rev. Lett. 129, 270502 \(2022\)](#).
- [51] Chi-Fang Chen, Hsin-Yuan Huang, Richard Kueng, and Joel A. Tropp, Concentration for random product formulas, [PRX Quantum 2, 040305 \(2021\)](#).
- [52] Etienne Granet and Henrik Dreyer, Hamiltonian dynamics on digital quantum computers without discretization error, [npj Quantum Inf. 10, 82 \(2024\)](#).
- [53] Chusei Kiumi and Bálint Koczor, TE-PAI: Exact time evolution by sampling random circuits, [Quantum Sci. Technol. 10, 045071 \(2025\)](#).
- [54] Sam McArdle, Earl Campbell, and Yuan Su, Exploiting fermion number in factorized decompositions of the electronic structure Hamiltonian, [Phys. Rev. A 105, 012403 \(2022\)](#).

- [55] Daniel Burgarth, Paolo Facchi, Alexander Hahn, Mattias Johnsson, and Kazuya Yuasa, Strong error bounds for Trotter & Strang-splittings and their implications for quantum chemistry, *Phys. Rev. Res.* **6**, 043155 (2024).
- [56] Craig Gidney, Noah Shutty, and Cody Jones, Magic state cultivation: growing T states as cheap as CNOT gates, [arXiv:2409.17595](https://arxiv.org/abs/2409.17595).
- [57] Earl Campbell, Shorter gate sequences for quantum computing by mixing unitaries, *Phys. Rev. A* **95**, 042306 (2017).
- [58] Vadym Kliuchnikov, Kristin Lauter, Romy Minko, Adam Paetznick, and Christophe Petit, Shorter quantum circuits via single-qubit gate approximation, *Quantum* **7**, 1208 (2023).
- [59] Bálint Koczor, John J. L. Morton, and Simon C. Benjamin, Probabilistic interpolation of quantum rotation angles, *Phys. Rev. Lett.* **132**, 130602 (2024).
- [60] Emiel Koridon, Saad Yalouz, Bruno Senjean, Francesco Buda, Thomas E. O'Brien, and Lucas Visscher, Orbital transformations to reduce the 1-norm of the electronic structure Hamiltonian for quantum computing applications, *Phys. Rev. Res.* **3**, 033127 (2021).
- [61] Ignacio Loaiza, Alireza Marefat Khah, Nathan Wiebe, and Artur F. Izmaylov, Reducing molecular electronic Hamiltonian simulation cost for linear combination of unitaries approaches, *Quantum Sci. Technol.* **8**, 035019 (2023).
- [62] Ignacio Loaiza and Artur F. Izmaylov, Block-invariant symmetry shift: Preprocessing technique for second-quantized Hamiltonians to improve their decompositions to linear combination of unitaries, *J. Chem. Theory Comput.* **19**, 8201 (2023).
- [63] Zhendong Li, Junhao Li, Nikesh S. Dattani, C. J. Umrigar, and Garnet Kin Chan, The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations, *J. Chem. Phys.* **150**, 024302 (2019).
- [64] Ian D. Kivlichan, Jarrod McClean, Nathan Wiebe, Craig Gidney, Alán Aspuru-Guzik, Garnet Kin-Lic Chan, and Ryan Babbush, Quantum simulation of electronic structure with linear depth and connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [65] Mario Motta, James Shee, Shiwei Zhang, and Garnet Kin-Lic Chan, Efficient ab initio auxiliary-field quantum Monte Carlo calculations in Gaussian bases via low-rank tensor decomposition, *J. Chem. Theory Comput.* **15**, 3510 (2019).
- [66] Bo Peng and Karol Kowalski, Highly efficient and scalable compound decomposition of two-electron integral tensor and its application in coupled cluster calculations, *J. Chem. Theory Comput.* **13**, 4179 (2017).
- [67] Daniel Litinski, Magic state distillation: Not as costly as you think, *Quantum* **3**, 205 (2019).
- [68] Yutaro Akahoshi, Kazunori Maruyama, Hirotaka Oshima, Shintaro Sato, and Keisuke Fujii, Partially fault-tolerant quantum computing architecture with error-corrected Clifford gates and space-time efficient analog rotations, *PRX Quantum* **5**, 010337 (2024).
- [69] Guang Hao Low, Vadym Kliuchnikov, and Nathan Wiebe, Well-conditioned multiproduct Hamiltonian simulation, [arXiv:1907.11679](https://arxiv.org/abs/1907.11679).
- [70] Almudena Carrera Vazquez, Daniel J. Egger, David Ochsner, and Stefan Woerner, Well-conditioned multiproduct formulas for hardware-friendly Hamiltonian simulation, *Quantum* **7**, 1067 (2023).
- [71] Gumaro Rendon, Jacob Watkins, and Nathan Wiebe, Improved accuracy for Trotter simulations using Chebyshev interpolation, *Quantum* **8**, 1266 (2024).
- [72] Junaid Aftab, Dong An, and Konstantina Trivisa, Multiproduct Hamiltonian simulation with explicit commutator scaling, [arXiv:2403.08922](https://arxiv.org/abs/2403.08922).
- [73] Maxine Luo and J. Ignacio Cirac, Efficient simulation of quantum chemistry problems in an enlarged basis set, [arXiv:2407.04432](https://arxiv.org/abs/2407.04432).
- [74] <https://doi.org/10.5281/zenodo.15387186>
- [75] Matthew Pocrnic, Matthew Hagan, Juan Carrasquilla, Dvira Segal, and Nathan Wiebe, Composite qDRIFT-product formulas for quantum and classical simulations in real and imaginary time, *Phys. Rev. Res.* **6**, 013224 (2024).
- [76] Dominic W. Berry and Howard M. Wiseman, Optimal states and almost optimal adaptive measurements for quantum interferometry, *Phys. Rev. Lett.* **85**, 5098 (2000).
- [77] Alicja Dutkiewicz, Stefano Polla, Maximilian Scheurer, Christian Gogolin, William J. Huggins, and Thomas E. O'Brien, Error mitigation and circuit division for early fault-tolerant quantum phase estimation, [arXiv:2410.05369](https://arxiv.org/abs/2410.05369).
- [78] Trygve Helgaker, Poul Jørgensen, and Jeppe Olsen, *Molecular Electronic-Structure Theory* (John Wiley & Sons, Ltd, Chichester, England, 2014), 1st ed.
- [79] Kosuke Mitarai, Kiichiro Toyozumi, and Wataru Mizukami, Perturbation theory with quantum signal processing, *Quantum* **7**, 1000 (2023).
- [80] William J. Huggins, Jarrod R. McClean, Nicholas C. Rubin, Zhang Jiang, Nathan Wiebe, K. Birgitta Whaley, and Ryan Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *npj Quantum Inf.* **7**, 23 (2021).
- [81] Jeffrey Cohn, Mario Motta, and Robert M. Parrish, Quantum filter diagonalization with compressed double-factorized Hamiltonians, *PRX Quantum* **2**, 040352 (2021).
- [82] Isaac H. Kim, Ye-Hua Liu, Sam Pallister, William Pol, Sam Roberts, and Eunseok Lee, Fault-tolerant resource estimate for quantum chemical simulations: Case study on Li-ion battery electrolyte molecules, *Phys. Rev. Res.* **4**, 023019 (2022).
- [83] Oumarou Oumarou, Maximilian Scheurer, Robert M. Parrish, Edward G. Hohenstein, and Christian Gogolin, Accelerating quantum computations of chemistry through regularized compressed double factorization, *Quantum* **8**, 1371 (2024).
- [84] Maximilian Mörchen, Guang Hao Low, Thomas Weymuth, Hongbin Liu, Matthias Troyer, and Markus Reiher, Classification of electronic structures and state preparation for quantum computation of reaction chemistry, [arXiv:2409.08910](https://arxiv.org/abs/2409.08910) [physics].
- [85] Athena Caesura, Christian L. Cortes, William Poll, Sukin Sim, Mark Steudtner, Gian-Luca R. Anselmetti, Matthias Degroote, Nikolaj Moll, Raffaele Santagati, Michael Streif, *et al.*, Faster quantum chemistry simulations on a quantum

- computer with improved tensor factorization and active volume compilation, [PRX Quantum 6, 030337 \(2025\)](#).
- [86] Sergey Bravyi, Jay M. Gambetta, Antonio Mezzacapo, and Kristan Temme, Tapering off qubits to simulate fermionic Hamiltonians, [arXiv:1701.08213](#).
- [87] Andrew Tranter, Peter J. Love, Florian Mintert, and Peter V. Coveney, A comparison of the Bravyi–Kitaev and Jordan–Wigner transformations for the quantum simulation of quantum chemistry, [J. Chem. Theory Comput. 14, 5617 \(2018\)](#).
- [88] Andrew Tranter, Peter J. Love, Florian Mintert, Nathan Wiebe, and Peter V. Coveney, Ordering of trotterization: Impact on errors in quantum simulation of electronic structure, [Entropy 21, 1218 \(2019\)](#).
- [89] Changhao Yi and Elizabeth Crosson, Spectral analysis of product formulas for quantum simulation, [npj Quantum Inf. 8, 37 \(2022\)](#).
- [90] Jakob Günther, Alberto Baiardi, Markus Reiher, and Matthias Christandl, More quantum chemistry with fewer qubits, [Phys. Rev. Res. 6, 043021 \(2024\)](#).
- [91] Alex Bocharov, Martin Roetteler, and Krysta M. Svore, Efficient synthesis of probabilistic quantum circuits with fallback, [Phys. Rev. A 91, 052317 \(2015\)](#).
- [92] Craig Gidney and Austin G. Fowler, Efficient magic state factories with a catalyzed $|CCZ\rangle \rightarrow 2|T\rangle$ transformation, [Quantum 3, 135 \(2019\)](#).
- [93] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, Improving quantum algorithms for quantum chemistry, [arXiv:1403.1539](#).
- [94] Andrew Tranter, Peter J. Love, Florian Mintert, and Peter V. Coveney, A comparison of the Bravyi–Kitaev and Jordan–Wigner transformations for the quantum simulation of quantum chemistry, [J. Chem. Theory Comput. 14, 5617 \(2018\)](#).
- [95] Earl T. Campbell, Early fault-tolerant simulations of the Hubbard model, [Quantum Sci. Technol. 7, 015007 \(2021\)](#).
- [96] Tzu-Ching Yen, Vladyslav Verteletskyi, and Artur F. Izmaylov, Measuring all compatible operators in one series of single-qubit measurements using unitary transformations, [J. Chem. Theory Comput. 16, 2400 \(2020\)](#).
- [97] Tzu-Ching Yen, Aadithya Ganeshram, and Artur F. Izmaylov, Deterministic improvements of quantum measurements with grouping of compatible operators, non-local transformations, and covariance estimates, [npj Quantum Inf. 9, 14 \(2023\)](#).