Nearly Optimal Lattice Simulation by Product Formulas

Andrew M. Childs and Yuan Su

Phys. Rev. Lett. **123**, 050503 — Published 2 August 2019

DOI: 10.1103/PhysRevLett.123.050503

Nearly optimal lattice simulation by product formulas

Andrew M. Childs and Yuan Su
Department of Computer Science, Institute for Advanced Computer Studies, and Joint Center for Quantum Information and Computer Science, University of Maryland
(Dated: July 8, 2019)

We consider simulating an n-qubit Hamiltonian with nearest-neighbor interactions evolving for time t on a quantum computer. We show that this simulation has gate complexity $(nt)^{1+o(1)}$ using product formulas, a straightforward approach that has been demonstrated by several experimental groups. While it is reasonable to expect this complexity—in particular, this was claimed without rigorous justification by Jordan, Lee, and Preskill—we are not aware of a straightforward proof. Our approach is based on an analysis of the local error structure of product formulas, as introduced by Descombes and Thalhammer and significantly simplified here. We prove error bounds for canonical product formulas, which include well-known constructions such as the Lie-Trotter-Suzuki formulas. We also develop a local error representation for time-dependent Hamiltonian simulation, and we discuss generalizations to periodic boundary conditions, constant-range interactions, and higher dimensions. Combined with a previous lower bound, our result implies that product formulas can simulate lattice Hamiltonians with nearly optimal gate complexity.

Simulating the Hamiltonian dynamics of a quantum system is one of the most natural applications of a quantum computer. Indeed, the idea of quantum computation, as suggested by Feynman [1] and others, was primarily motivated by the problem of quantum simulation. Quantum computers can simulate a variety of physical systems, including quantum chemistry [2–5], quantum field theory [6, 7], and many-body physics [8], and could ultimately lead to practical applications such as designing new pharmaceuticals, catalysts, and materials [9, 10].

A natural class of Hamiltonians that includes many physically reasonable systems is the class of lattice Hamiltonians [6, 11–13]. Lattice Hamiltonians arise in many models of condensed matter physics, including systems of spins (e.g., Ising, XY, and Heisenberg models; Kitaev's toric code and honeycomb models; etc.), fermions (e.g., the Hubbard model and the t-J model), and bosons (e.g., the Bose-Hubbard model). Note that fermion models can be simulated using local interactions among qubits by using a mapping to qubits that preserves locality [14]. Digital simulations of quantum field theory also typically involve approximation by a lattice system [6].

For simplicity, we mainly focus on nearest-neighbor lattice systems in one dimension (although we discuss generalizations to other lattice models as well). In this case, n qubits are laid out on a one-dimensional lattice and the Hamiltonian only involves nearest-neighbor interactions. Specifically, a Hamiltonian H is a lattice Hamiltonian if it acts on n qubits and can be decomposed as $H = \sum_{j=1}^{n-1} H_{j,j+1}$, where each $H_{j,j+1}$ is a Hermitian operator that acts nontrivially only on qubits j and $j+1$. We assume that $\max_j \|H_{j,j+1}\| \leq 1$, for otherwise we evolve under the normalized Hamiltonian $H/\max_j \|H_{j,j+1}\|$ for time $\max_j \|H_{j,j+1}\| t$.

Lloyd's original proposal for an explicit quantum simulation algorithm [15] uses the Lie-Trotter product formula. Subsequent work achieves better asymptotic complexity [16] using higher-order Suzuki formulas [17]. We refer to all such formulas as *product formulas*. The product-formula algorithm is straightforward yet surprisingly efficient for quantum simulation. Indeed, it can conserve certain symmetries of the dynamics [18] and appears to be advantageous for various practical systems [2, 19, 20]. Although recent simulation algorithms have better asymptotic complexities [21–30], the product-formula approach remains a natural choice for experimental simulations [31–33] due to its simplicity and the fact that it does not require any ancilla qubits. Its study has also illuminated areas beyond quantum computing [34].

One of the main challenges in quantum simulation is to analyze the gate complexity of simulation algorithms. Explicit gate counts are especially desirable for near-term simulation because early quantum computers will only be able to reliably perform a limited number of gates. While existing analysis appears to be tight for recent simulation algorithms, the product-formula bound can be loose by several orders of magnitude [2, 8, 19, 20]. This dramatic gap makes it hard to identify the fastest simulation algorithm and to find optimized implementations for near-term applications [19].

Product formulas can simulate a lattice system with fixed accuracy with gate complexity $O\big(n(nt)^2\big)$ in the first-order case and $O\big(n(nt)^{1+\frac{1}{2k}}\big)$ in the

$(2k)$th-order case. However, it is natural to expect a more efficient simulation. Roughly speaking, a system simulates its own evolution for constant time using only constant circuit depth—and hence an extensive number of gates—so one might expect a true simulation complexity of $O(nt)$. Indeed, Jordan, Lee, and Preskill claimed that product formulas can simulate an n-qubit lattice system with $(nt)^{1+o(1)}$ gates [6], but they did not provide rigorous justification and it is unclear how to formalize their argument. Subsequent work improves the analysis of the product-formula algorithm using information about commutation among terms in the Hamiltonian [3, 4, 19, 35], the distribution of norms of terms [36], and by randomizing the ordering of terms [37, 38]. However, none of these improvements can achieve the claimed gate complexity $(nt)^{1+o(1)}$ for lattice simulation.

Main result. Let $H = \sum_{j=1}^{n-1} H_{j,j+1}$ be an n-qubit lattice Hamiltonian. We order the terms in the *even-odd* pattern $H_{1,2}, H_{3,4}, \ldots, H_{2,3}, H_{4,5}, \ldots$ obtaining the first-order product formula

$$\begin{aligned} \mathscr{S}_1(t) :&= \prod_{k=1}^{\frac{n}{2}-1} e^{-itH_{2k,2k+1}} \prod_{k=1}^{\frac{n}{2}} e^{-itH_{2k-1,2k}} \\ &= e^{-itH_{\text{even}}} e^{-itH_{\text{odd}}} \end{aligned} \tag{1}$$

and the $(2k)$th-order product formulas

$$\begin{aligned} \mathscr{S}_2(t) &:= e^{-i\frac{t}{2}H_{\text{odd}}} e^{-itH_{\text{even}}} e^{-i\frac{t}{2}H_{\text{odd}}} \\ \mathscr{S}_{2k}(t) &:= \mathscr{S}_{2k-2}(p_k t)^2 \, \mathscr{S}_{2k-2}((1-4p_k)t) \, \mathscr{S}_{2k-2}(p_k t)^2 \end{aligned} \tag{2}$$

with $p_k := 1/(4 - 4^{1/(2k-1)})$. Our main result is an asymptotic upper bound on the product-formula error

$$\begin{aligned} \left\| \mathscr{S}_1(t) - e^{-itH} \right\| &= O(nt^2) \\ \left\| \mathscr{S}_{2k}(t) - e^{-itH} \right\| &= O(nt^{2k+1}), \end{aligned} \tag{3}$$

where $\|\cdot\|$ denotes the spectral norm.

The above error bound works well only for very small t. To simulate for a longer time, we divide the entire evolution into r segments, and within each segment, we simulate using product formulas. To achieve accuracy ϵ, it suffices to choose $r_1 = O(nt^2/\epsilon)$ for the first-order formula and $r_{2k} = O\big(t(nt/\epsilon)^{\frac{1}{2k}}\big)$ for the $(2k)$th-order formula. Equivalently, we have gate complexity $O\big((nt)^2\big)$ and $O\big((nt)^{1+\frac{1}{2k}}\big)$ for the first- and $(2k)$th-order algorithm, assuming that we simulate with constant accuracy.

For any $\delta > 0$, we choose an integer k sufficiently large so that $\frac{1}{2k} \le \delta$, upper-bounding the gate complexity as $O\big((nt)^{1+\frac{1}{2k}}\big) = O\big((nt)^{1+\delta}\big)$. This proves that the product-formula algorithm has asymptotic gate complexity $(nt)^{1+o(1)}$. Combining with the lower bound of $\widetilde{\Omega}(nt)$ established in [12], we have showed that product formulas can simulate a lattice Hamiltonian with nearly optimal gate complexity.

Applications. As an immediate application, our result gives a rigorous proof of the Jordan-Lee-Preskill claim about the complexity of simulating quantum field theory [6]. Recent works have analyzed the gate complexity of other quantum field theory simulations [39], including digital simulation of gauge theories [40]. The lattice Hamiltonians there have similar locality, so our analysis still applies. We expect our technique can be generalized to speed up the simulation of other systems, such as electronic structure Hamiltonians [9], power-law decaying interactions [41], exponentially decaying interactions [42], and clustered Hamiltonians [43], but we leave a thorough study of such generalizations as a subject for future work [44].

To simulate an n-qubit lattice Hamiltonian for time t, our algorithm has circuit depth $n^{o(1)}t^{1+o(1)}$. As a side application, our analysis gives a tensor network representation of lattice systems with bond dimension $2^{n^{o(1)}t^{1+o(1)}}$, using the counting argument of [45]. This exponentially improves a recent construction of [46, Lemma 17] which uses only the first-order Trotter decomposition.

We work primarily with an idealized setting where quantum operations can be performed faithfully. However, in realistic experiments, quantum gates will be imperfectly implemented. For such a case, Reference [47] estimates the simulation accuracy as $\frac{\alpha}{r^{2k}} + \beta r$ in diamond-norm distance [48, 49], where α captures the algorithmic error of product formulas and β captures gate errors. This leads to an optimal number of segments $r = \big(\alpha 2k/\beta\big)^{\frac{1}{2k+1}}$, which can be improved by our result. Specifically, the original analysis in [16] implies $\alpha_{\text{orig}} = O\big((nt)^{2k+1}\big)$. This has been improved by subsequent work [19, 37], although none of these improvements achieves linear scaling in n. In contrast, the analysis of this letter gives $\alpha_{\text{opt}} = O\big(nt^{2k+1}\big)$, improving the performance as a function of n even in the presence of noise.

Our main goal is to establish the gate complexity of $(nt)^{1+o(1)}$ for the product-formula algorithm. However, our analysis is not only nearly optimal in the asymptotic regime but also appears to be much tighter in practice. For concreteness, we numerically implement and optimize our fourth-order bound, and compare it with previous product-formula analysis, for simulation of a one-dimensional Heisenberg model with a random magnetic field with open

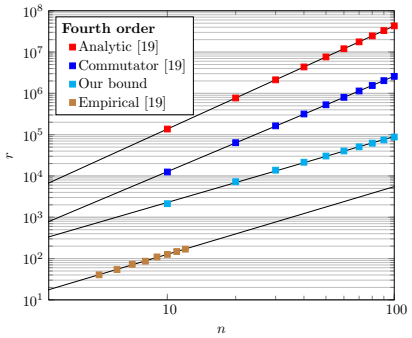

FIG. 1: Comparison of r for different product-formula bounds for the Heisenberg model (see [50, Section VI] for detailed parameters). Error bars are omitted as they are negligibly small on the plot. Straight lines show power-law fits to the data.

boundary conditions [50, Eq.(98)] (see Figure 1). We find that the scaling of our bound matches the empirical performance and the constant prefactor is off by only one order of magnitude, a significant improvement over previous rigorous bounds [19]. Further improvements of our bound are possible by optimizing its numerical implementation; we leave a detailed study for future work [44].

Analysis of the first-order algorithm. The key technique behind our approach is an integral representation of the error $\mathscr{S}(t) - e^{-itH}$ that we develop based on Descombes and Thalhammer's *local error analysis* of product formulas [51]. In the local error representation, the integrand is expressed as a linear combination of nested commutators, where the numbers of summands and nesting layers are both independent of n and t. We use this representation to get the correct asymptotic gate count as a function of n and t. In contrast, the conventional approach uses the Baker-Campbell-Hausdorff formula or naive Taylor expansion, which requires the manipulation of infinite series and appears to be technically challenging to analyze [52] [50, Section I].

To illustrate the proof idea, we show how to obtain $\left\|\mathscr{S}_1(t) - e^{-itH}\right\| = O(nt^2)$ for the first-order formula. We differentiate $\mathscr{S}_1(t)$ and obtain

$$\mathscr{S}_1'(t) = -iH\mathscr{S}_1(t) + \left[e^{-itH_{\text{even}}},\ -iH_{\text{odd}}\right]e^{-itH_{\text{odd}}}. \tag{4}$$

Using the variation-of-constants formula [51] [53, Theorem 4.9] with initial condition $\mathscr{S}_1(0) = I$, we find an integral representation of the product-formula error $\mathscr{S}_1(t) - e^{-itH}$ as

$$\int_0^t \mathrm{d}\tau_1\ e^{-i(t-\tau_1)H}\left[e^{-i\tau_1 H_{\text{even}}},\ -iH_{\text{odd}}\right]e^{-i\tau_1 H_{\text{odd}}}. \tag{5}$$

We repeat this procedure to analyze the commutator $\left[e^{-i\tau_1 H_{\text{even}}},\ -iH_{\text{odd}}\right]$, obtaining an upper bound on the spectral-norm error

$$\left\|\mathscr{S}_1(t) - e^{-itH}\right\| \le \int_0^t \mathrm{d}\tau_1 \int_0^{\tau_1} \mathrm{d}\tau_2 \left\|\left[H_{\text{even}},\ H_{\text{odd}}\right]\right\|. \tag{6}$$

We expand H_{odd} and H_{even} according to their definitions. Fixing an arbitrary term $H_{2k-1,2k}$ in H_{odd}, the commutator $\left[H_{2l,2l+1},\ H_{2k-1,2k}\right]$ is nonzero only when $l \in \{k-1, k\}$. We thus find that

$$\left[H_{\text{even}},\ H_{\text{odd}}\right] = \sum_{k=1}^{\frac{n}{2}} \left[H_{2k-2,2k-1} + H_{2k,2k+1},\ H_{2k-1,2k}\right]. \tag{7}$$

Using the triangle inequality, we have $\left\|\mathscr{S}_1(t) - e^{-itH}\right\| = O(nt^2)$, which proves the claim (3) for the first-order case.

Ordering robustness. Our above bound works when terms of the lattice Hamiltonian are ordered in the even-odd pattern. However, this choice is not necessary: the first-order algorithm has the same asymptotic error bound with respect to any ordering of the lattice terms.

Our analysis relies on an error bound for swapping lattice terms:

$$\left\|\left[e^{-itH_{k,k+1}}, e^{-itH_{l,l+1}}\right]\right\| \le 2t^2 \tag{8}$$

if $|k-l| = 1$ and 0 otherwise. In words, we may swap two exponentials $e^{-itH_{k,k+1}}$ and $e^{-itH_{l,l+1}}$ without error unless their supports overlap, in which case the error is $O(t^2)$.

Let $H = \sum_{j=1}^{n-1} H_{j,j+1}$ be a lattice Hamiltonian. We now simulate it using the first-order formula, but allow terms to be ordered arbitrarily as $\prod_{j=1}^{n-1} e^{-itH_{\sigma(j),\sigma(j)+1}}$, where $\sigma \in S_{n-1}$ is a permutation on the $n-1$ elements $\{1, \ldots, n-1\}$. Then the spectral-norm error $\left\|\prod_{j=1}^{n-1} e^{-itH_{\sigma(j),\sigma(j)+1}} - e^{-itH}\right\|$ is upper bounded by

$$\left\|\prod_{j=1}^{n-1} e^{-itH_{\sigma(j),\sigma(j)+1}} - e^{-itH_{\text{even}}}e^{-itH_{\text{odd}}}\right\| + \left\|e^{-itH_{\text{even}}}e^{-itH_{\text{odd}}} - e^{-itH}\right\|. \tag{9}$$

The second term is upper bounded by $O(nt^2)$. For the first term, we transform $\prod_{j=1}^{n-1} e^{-itH_{\sigma(j),\sigma(j)+1}}$ into $\prod_{k=1}^{\frac{n}{2}-1} e^{-itH_{2k,2k+1}} \prod_{k=1}^{\frac{n}{2}} e^{-itH_{2k-1,2k}}$ by swapping neighboring exponentials. Every time two exponentials are swapped, we use (8) to bound the error. The total number of swaps of exponentials $e^{-itH_{k,k+1}}$ and $e^{-itH_{l,l+1}}$ with $|k-l| = 1$ is at most $2n$, incurring error $4nt^2 = O(nt^2)$.

We have therefore obtained the same asymptotic error bound for an arbitrary ordering of the Hamiltonian terms. We call this phenomenon *ordering robustness*. Our analysis shows that the first-order algorithm is ordering-robust. Whether a similar property holds for a general higher-order formula remains an open question.

We also numerically compare the first-order algorithm with the even-odd ordering and the ordering of [19]. Although they have the same asymptotic error bound, in practice the even-odd ordering has smaller exponent and constant prefactor. Details can be found in [50, Section VI].

Analysis of higher-order algorithms. Analyzing higher-order product formulas is more challenging. To this end, we represent them in a *canonical form*, which is easy to manipulate and encompasses well-known constructions such as the Lie-Trotter-Suzuki formulas $\mathscr{S}_1(t)$, $\mathscr{S}_{2k}(t)$ as special cases. We then use the variation-of-constants formula to write

$$\mathscr{S}_{2k}(t) - e^{-itH} = \int_0^t e^{-i(t-\tau)H} \mathscr{S}_{2k}(\tau) \mathscr{T}(\tau) \mathrm{d}\tau, \tag{10}$$

where $\mathscr{T}(t) = \mathscr{S}_{2k}^\dagger(t)\big[\frac{\mathrm{d}}{\mathrm{d}t}\mathscr{S}_{2k}(t) - (-iH)\mathscr{S}_{2k}(t)\big]$. As a $(2k)$th-order formula, $\mathscr{S}_{2k}(t)$ satisfies an *order condition* $\mathscr{S}_{2k}(t) = e^{-itH} + O(t^{2k+1})$, which further implies by Taylor's theorem

$$\mathscr{T}(\tau) = 2k \int_0^1 \mathrm{d}x\ (1-x)^{2k-1} \mathscr{T}^{(2k)}(x\tau) \frac{\tau^{2k}}{(2k)!}. \tag{11}$$

Canonical product formulas and their order conditions are further discussed in [50, Section II].

A direct expansion of $\mathscr{T}(t)$ gives the correct t-dependence $O(t^{2k+1})$ of the product-formula error, but the scaling in n is incorrect. Instead, we seek an alternative expression for the integrand that consists of a linear combination of nested commutators, where the number of commutators and nested layers are both independent of n and t. Such an expression is referred to as a *local error representation* in [51]. However, the result of [51] depends on auxiliary functions whose recursive structure is hard to unravel. Instead, we develop a simpler representation of the local error structure [50, Section III].

In our representation, the operator $\mathscr{T}(\tau)$ can be written as a linear combination of operator-valued functions of the form $e^{i\tau X_1} \cdots e^{i\tau X_l} Y e^{-i\tau X_l} \cdots e^{-i\tau X_1}$, where operators $X_j, Y \in \{H_{\mathrm{even}}, H_{\mathrm{odd}}\}$. As such, its higher-order derivatives consist of unitary conjugations and commutators. When a commutator is composed, we imitate (7) to show that the support of the operator is expanded by at most a constant factor. When a unitary conjugation is composed, we decompose the unitary operators and cancel exponentials with non-overlapping supports. Throughout this procedure, we only introduce $O(n)$ error in the innermost layer, proving the claim in (3) for the higher-order cases. This error analysis is discussed in more details in [50, Section IV].

Generalized lattice Hamiltonians. We have so far focused on time-independent one-dimensional systems with nearest-neighbor interactions and open boundary conditions. However, our analysis can be easily adapted to handle time-dependent Hamiltonians, periodic boundary conditions, constant-range interactions, and higher-dimensional systems, again with nearly optimal gate complexity.

When the Hamiltonian $H(t)$ is time-dependent, the problem of quantum simulation becomes more difficult [54]. Then there no longer exists a closed-form solution to the Schrödinger equation. Furthermore, some quantum simulation algorithms [22, 25] that behave well in the time-independent case fail to handle the time-dependent Hamiltonian simulation. Nevertheless, we show that product formulas can simulate time-dependent lattice Hamiltonians with nearly optimal gate complexity. We group the terms in the even-odd pattern

$$\begin{aligned} H_{\mathrm{odd}}(t) &= H_{1,2}(t) + H_{3,4}(t) + \cdots \\ H_{\mathrm{even}}(t) &= H_{2,3}(t) + H_{4,5}(t) + \cdots \end{aligned} \tag{12}$$

and simulate using the time-dependent Lie-Trotter-Suzuki formulas $\mathscr{S}_{\mathcal{T},2k}(t)$ [54]. We show that

$$\left\| \mathscr{S}_{\mathcal{T},2k}(t) - \exp_{\mathcal{T}}\left(-i\int_0^t \mathrm{d}v\ H(v)\right) \right\| = O\big(nt^{2k+1}\big) \tag{13}$$

where $\exp_{\mathcal{T}}$ denotes the time-ordered matrix exponential. Similar to the time-independent case, we find that the total gate complexity is $O\big((nt)^{1+\frac{1}{2k}}\big)$. See [50, Section V] for detailed discussions.

We also consider lattice Hamiltonians with periodic boundary conditions $H = \sum_{j=1}^{n-1} H_{j,j+1} + H_{1,n}$, where $H_{j,k}$ represents a local term that acts nontrivially only on qubits j and k. To simulate such a system, we decompose H as $H = H_{\mathrm{bndry}} + H_{\mathrm{even}} + H_{\mathrm{odd}}$, where $H_{\mathrm{bndry}} = H_{1,n}$. Correspondingly, we also use a canonical product formula with three exponentials per stage. With a similar analysis for the open boundary condition, we find that the product-formula error is $O(nt^{2k+1})$ as expected.

A generalization of this approach allows us to simulate a D-dimensional lattice Hamiltonian with nearly optimal gate complexity. We use a $2D$-coloring of the edges of the lattice to decompose the

Hamiltonian into $2D$ terms, each of which is a sum of commuting terms. We also extend the definition of canonical product formulas to allow for $2D$ exponentials per stage. An analysis of the local error structure shows that this algorithm has gate complexity $O((L^D t)^{1+\frac{1}{2k}}/\epsilon^{\frac{1}{2k}}) = O((nt)^{1+\frac{1}{2k}}/\epsilon^{\frac{1}{2k}})$, where n is the total number of lattice sites and $L = n^{\frac{1}{D}}$ is the linear size of the lattice.

Finally, our algorithm can also simulate lattice Hamiltonians with constant-range interactions $H = \sum_{j=1}^{n-\ell+1} H_{j,\ldots,j+\ell-1}$. To achieve nearly-optimal gate complexity, we classify the Hamiltonian terms into the ℓ groups

$$\begin{aligned} H_{[1]} &= H_{1,\ldots,\ell} + H_{\ell+1,\ldots,2\ell} + \cdots \\ H_{[2]} &= H_{2,\ldots,\ell+1} + H_{\ell+2,\ldots,2\ell+1} + \cdots \\ &\vdots \\ H_{[\ell]} &= H_{\ell,\ldots,2\ell-1} + H_{2\ell,\ldots,3\ell-1} + \cdots \end{aligned} \tag{14}$$

and use a product formula with ℓ elementary exponentials per stage.

Discussion. The product-formula algorithm is arguably the simplest approach to quantum simulation. We have showed that this approach can simulate lattice Hamiltonians with nearly optimal gate complexity. Our algorithm invokes product formulas by ordering terms in an even-odd pattern, which is conceptually easy to understand and straightforward to implement. Beyond the one-dimensional time-independent system with nearest-neighbor interactions and open boundary conditions, our analysis is also applicable to periodic boundary conditions, constant-range interactions, higher dimensions, and the time-dependent case, all with nearly optimal gate complexity. Our result also gives product-formula bounds that are much tighter in practice.

Recently, Haah, Hastings, Kothari and Low (HHKL) proposed another nearly optimal algorithm for lattice simulation [12]. Instead of analyzing the product-formula approach, they develop a new approach motivated by the Lieb-Robinson bound [55, 56], which quantifies how fast information can propagate in a system with local interactions. HHKL decomposes the entire evolution into blocks, where each block involves forward and backward evolution on a small region. Using product formulas within each block, their approach gives an ancilla-free algorithm for lattice simulation with asymptotic gate complexity $(nt)^{1+o(1)}$. However, this results in a much larger constant prefactor in practice than the pure product-formula algorithm analyzed here [50, Section VI].

The near optimality of HHKL depends essentially on the use of a Lieb-Robinson bound. As noted in [12], it may be difficult to apply this idea to Hamiltonians whose interactions are described by general graphs. Our approach directly exploits the commutation of lattice terms without the help of Lieb-Robinson bounds, which we expect could illuminate the simulation of other physical systems [9, 41–43].

Our local error analysis represents the product-formula error as an integral of a linear combination of nested commutators. Similar techniques have been used to establish the Lieb-Robinson bound and to study computational complexity aspects of many-body physics [41, 55–57]. We leave it as an avenue for future work to explore whether these techniques could find more applications in the study of locality in quantum systems.

ACKNOWLEDGMENTS

Y.S. thanks Nathan Wiebe, Guang Hao Low, Minh Cong Tran, Stephen Jordan, Jeongwan Haah, Robin Kothari, Su-Kuan Chu, Rolando Somma, Xin Wang, Alexey Gorshkov, James R. Garrison, Anurag Anshu, Xiaodi Wu, Leonard Wossnig, Salini Karuvade, Xiao Yuan, Simon Benjamin, Shuchen Zhu, Scott Lawrence, and Zohreh Davoudi for helpful discussions. We thank anonymous referees for their comments. This work was supported in part by the Army Research Office (MURI award W911NF-16-1-0349), the Canadian Institute for Advanced Research, the National Science Foundation (grants 1526380 and 1813814), and the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Teams and Quantum Testbed Pathfinder programs.

[1] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).
[2] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, Physical Review A **91**, 022311 (2015), arXiv:1410.8159.
[3] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, Physical Review A **90**, 022305 (2014), arXiv:1312.1695.
[4] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, Quantum Information and Computation **15**, 361 (2015), arXiv:1406.4920.
[5] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and

A. Aspuru-Guzik, "Quantum chemistry in the age of quantum computing," (2018), arXiv:1812.09976.
[6] S. P. Jordan, K. S. M. Lee, and J. Preskill, Science **336**, 1130 (2012), arXiv:1111.3633.
[7] S. P. Jordan, K. S. M. Lee, and J. Preskill, Quantum Information and Computation **14**, 1014 (2014), arXiv:1112.4833.
[8] S. Raeisi, N. Wiebe, and B. C. Sanders, New Journal of Physics **14**, 103017 (2012), arXiv:1108.4318.
[9] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven, and G. K.-L. Chan, Physical Review X **8**, 011044 (2018), arXiv:1706.00023.
[10] R. de Wolf, Ethics and Information Technology **19**, 271 (2017), arXiv:1712.05380.
[11] D. J. Luitz, N. Laflorencie, and F. Alet, Physical Review B **91**, 081103 (2015), arXiv:1411.0660.
[12] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, "Quantum algorithm for simulating real time evolution of lattice Hamiltonians," (2018), arXiv:1801.03922.
[13] X. Gao, S.-T. Wang, and L.-M. Duan, Physical Review Letters **118**, 040502 (2017), arXiv:1607.04947.
[14] F. Verstraete and J. I. Cirac, Journal of Statistical Mechanics **2005**, P09012 (2005), arXiv:cond-mat/0508353.
[15] S. Lloyd, Science **273**, 1073 (1996).
[16] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Communications in Mathematical Physics **270**, 359 (2007), arXiv:quant-ph/0508139.
[17] M. Suzuki, Journal of Mathematical Physics **32**, 400 (1991).
[18] N. Hatano and M. Suzuki, in *Quantum Annealing and Other Optimization Methods*, edited by A. Das and B. K. Chakrabarti (Springer, 2005) pp. 37–68.
[19] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Proceedings of the National Academy of Sciences **115**, 9456 (2018), arXiv:1711.10980.
[20] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Proceedings of the National Academy of Sciences **114**, 7555 (2017), arXiv:1605.03590.
[21] D. W. Berry, A. M. Childs, and R. Kothari, in *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science* (2015) pp. 792–809, arXiv:1501.01715.
[22] D. W. Berry and A. M. Childs, Quantum Information and Computation **12**, 29 (2012), arXiv:0910.4157.
[23] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Physical Review Letters **114**, 090502 (2015), arXiv:1412.4687.
[24] G. H. Low and I. L. Chuang, "Hamiltonian simulation by qubitization," (2016), arXiv:1610.06546.
[25] G. H. Low and I. L. Chuang, Physical Review Letters **118**, 010501 (2017), arXiv:1606.02685.
[26] G. H. Low, "Hamiltonian simulation with nearly optimal dependence on spectral norm," (2018), arXiv:1807.03967.
[27] D. W. Berry, R. Cleve, and S. Gharibian, Quantum Information & Computation **14**, 1 (2014), arXiv:1211.4637.
[28] E. Campbell, "A random compiler for fast Hamiltonian simulation," (2018), arXiv:1811.08017.
[29] M. Kieferová, A. Scherer, and D. Berry, Physical Review A **99**, 042314 (2019), arXiv:1805.00582.
[30] G. H. Low and N. Wiebe, "Hamiltonian simulation in the interaction picture," (2018), arXiv:1805.00675.
[31] K. R. Brown, R. J. Clark, and I. L. Chuang, Physical Review Letters **97**, 050504 (2006), arXiv:quant-ph/0601021.
[32] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, E. Solano, and J. M. Martinis, Nature Communications **6**, 7654 (2015), arXiv:1501.07703.
[33] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos, Science **334**, 57 (2011), arXiv:1109.1512.
[34] S. Blanes and F. Casas, *A Concise Introduction to Geometric Numerical Integration* (Chapman and Hall/CRC, 2016).
[35] R. D. Somma, Journal of Mathematical Physics **57**, 062202 (2016), arXiv:1512.03416.
[36] S. Hadfield and A. Papageorgiou, New Journal of Physics **20**, 043003 (2018).
[37] A. M. Childs, A. Ostrander, and Y. Su, "Faster quantum simulation by randomization," (2018), arXiv:1805.08385.
[38] C. Zhang, in *Monte Carlo and Quasi-Monte Carlo Methods 2010*, edited by L. Plaskota and H. Woźniakowski (Springer, 2012) pp. 709–719.
[39] J. Preskill, "Simulating quantum field theory with a quantum computer," (2018), arXiv:1811.10085.
[40] H. Lamm, S. Lawrence, and Y. Yamauchi, "Quantum simulation of gauge theories," (2019), arXiv:1903.08807.
[41] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, to appear in Physical Review X (2019), arXiv:1808.05225.
[42] R. Machleidt and D. R. Entem, Physics Reports **503**, 1 (2011), arXiv:1105.2919.
[43] T. Peng, A. Harrow, M. Ozols, and X. Wu, "Simulating large quantum circuits on a small quantum computer," (2019), arXiv:1904.00102.
[44] A. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, (2019), manuscript in preparation.
[45] R. Jozsa, "On the simulation of quantum circuits," (2006), arXiv:quant-ph/0603163.
[46] M. Holzäpfel and M. B. Plenio, "Efficient certification and simulation of local quantum many-body Hamiltonians," (2017), arXiv:1712.04396.
[47] G. C. Knee and W. J. Munro, Physical Review A **91**, 052327 (2015), arXiv:1502.04536.
[48] J. Watrous, *The theory of quantum information* (Cambridge University Press, 2018).
[49] M. M. Wilde, *Quantum Information Theory* (Cam-

bridge University Press, 2017).
[50] See Supplemental Material for details, which includes Refs. [58–61].
[51] S. Descombes and M. Thalhammer, BIT Numerical Mathematics **50**, 729 (2010).
[52] J. Haah and S. Jordan, Private communication (2018).
[53] A. W. Knapp, *Basic real analysis* (Birkhëuser, 2005).
[54] N. Wiebe, D. Berry, P. Høyer, and B. C. Sanders, Journal of Physics A **43**, 065203 (2010), arXiv:0812.0562.
[55] T. J. Osborne, Physical Review Letters **97**, 157202 (2006), arXiv:quant-ph/0508031.
[56] M. B. Hastings, "Locality in quantum systems," (2010), arXiv:1008.5137.
[57] D. Aharonov, I. Arad, U. Vazirani, and Z. Landau, New Journal of Physics **13**, 113043 (2011), arXiv:1011.3445.
[58] W. Auzinger and W. Herfort, Opuscula Mathematica **34**, 243 (2014).
[59] W. Rossmann, *Lie groups: An introduction through linear groups* (Oxford University Press, 2006).
[60] S. Chakraborty, A. Gilyén, and S. Jeffery, "The power of block-encoded matrix powers: improved regression techniques via faster Hamiltonian simulation," (2018), arXiv:1804.01973.
[61] R. Bellman, *Introduction to matrix analysis* (Society for Industrial and Applied Mathematics, 1997).