

On the complexity of implementing Trotter steps

Guang Hao Low,¹ Yuan Su,¹ Yu Tong,^{2,3} and Minh C. Tran^{4,5}

Abstract

Quantum dynamics can be simulated on a quantum computer by exponentiating elementary terms from the Hamiltonian in a sequential manner. However, such an implementation of Trotter steps has gate complexity depending on the total Hamiltonian term number, comparing unfavorably to algorithms using more advanced techniques.

We develop methods to perform faster Trotter steps with gate complexity sublinear in the number of terms. We focus on a class of 2-local Hamiltonians in one spatial dimension whose interaction strength decays with distance x according to power law $1/x^\alpha$ ($\alpha > 0$), although we discuss extensions to higher spatial dimensions as well. Naively, Trotter steps for n -qubit power-law interactions require $\Omega(n^2)$ gates to implement.

Our first method is based on block encodings of power-law systems with efficiently computable coefficients. While block encodings typically slow down quantum simulation by a factor proportional to the 1-norm of Hamiltonian coefficients, we overcome this barrier using a recursive Trotter decomposition to effectively reduce the norm. We obtain further improvements by simulating commuting Hamiltonian terms with an average combination cost. The resulting complexity is almost linear in the spacetime volume for $\alpha \geq 2$ and improves the state of the art for $\alpha < 2$.

We also show that Trotter steps can be implemented more efficiently when certain blocks of Hamiltonian coefficients exhibit low-rank properties. In particular, using a recursive low-rank decomposition, we show that power-law Hamiltonians can be simulated with gate complexity nearly linear in the spacetime volume for all $\alpha \geq 1$.

We apply our methods to simulate electronic structure Hamiltonians in second quantization in real space. Combining with a tighter error analysis, we show that it suffices to use $\left(\eta^{1/3}n^{1/3} + \frac{n^{2/3}}{\eta^{2/3}}\right)n^{1+o(1)}$ gates to simulate uniform electron gas in real space with n spin orbitals and η electrons, asymptotically improving the best results from previous work. We obtain an analogous result when the external potential of nuclei is introduced under the Born-Oppenheimer approximation.

We prove a circuit lower bound when the Hamiltonian coefficients take a continuum range of values. Specifically, we construct a class of 2-local Hamiltonians with commuting terms that requires at least $\Omega(n^2)$ gates to evolve with accuracy $\epsilon = \Omega(1/\text{poly}(n))$ for time $t = \Omega(\epsilon)$. Our proof is based on a gate-efficient reduction from the approximate synthesis of diagonal unitaries within the Hamming weight-2 subspace, which may be of independent interest.

Our result thus suggests the use of Hamiltonian structural properties as both necessary and sufficient to implement Trotter steps with lower gate complexity.

¹Microsoft Quantum, Redmond, WA 98052, USA

²Department of Mathematics, University of California, Berkeley, CA 94720, USA

³Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA 91125, USA

⁴Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁵Department of Physics, Harvard University, Cambridge, MA 02138, USA

Contents

1	Introduction	3
1.1	Quantum algorithms for quantum simulation	3
1.2	Previous related work	5
1.3	Faster Trotter steps	6
1.4	Circuit lower bound	9
2	Preliminaries	10
2.1	Notation and terminology	10
2.2	Product formulas	12
2.3	Block encoding and qubitization	14
2.4	Volume of diagonal unitaries	16
3	Faster Trotter steps using block encoding	17
3.1	Reducing the 1-norm of Hamiltonian coefficients	18
3.2	Preparation and selection subroutines	21
3.3	Summary of the algorithm	25
4	Faster Trotter steps using average-cost simulation	26
4.1	Simulating commuting terms with average combination cost	26
4.2	Preparation and selection subroutines	29
4.3	Summary of the algorithm	31
5	Faster Trotter steps using low-rank decomposition	32
5.1	Recursive low-rank decomposition	32
5.2	Summary of the algorithm	36
6	Applications to real-space quantum simulation	36
7	First circuit lower bound	39
7.1	Approximate synthesis of diagonal unitaries	40
7.2	Reduction from simulation in the Hamming weight-2 subspace	42
8	Discussion	44
A	Second circuit lower bound	46
A.1	Approximate synthesis of discrete diagonal unitaries	46
A.2	Reduction from the coefficient oracle	48
B	Trotter error with fermionic induced 1-norm scaling	49
C	Generalization	52
C.1	Higher spatial dimensionality	52
C.2	κ -local Hamiltonians	54
C.3	General fermionic Hamiltonians	55

1 Introduction

Many-body Hamiltonians can be efficiently simulated on digital quantum computers using either product formulas (such as the Lie-Trotter-Suzuki formulas) or more advanced simulation algorithms. While short steps of product formulas (known as *Trotter steps*) typically require an implementation cost proportional to the number of Hamiltonian terms, a host of techniques have been developed in recent years to implement other quantum simulation algorithms with significantly reduced complexities.

The purpose of this work is to develop methods for performing Trotter steps that go beyond the naive circuit implementation mentioned above, leading to faster quantum simulation algorithms. While our methods in principle apply to more general systems, we focus on 2-local Hamiltonians with power-law decaying interactions—those that decay with their distance x according to a power law $1/x^\alpha$ ($\alpha > 0$)—for concreteness. Many systems of physical relevance can be modeled by power-law interactions, such as trapped ions, Rydberg atoms, ultracold atoms and molecules, nitrogen-vacancy centers, and superconducting systems.

We first develop a block-encoding-based method to simulate power-law Hamiltonians with efficiently computable coefficients. While a block encoding introduces a slow-down factor proportional to the 1-norm of Hamiltonian coefficients which is large for power-law interactions, we overcome this barrier by recursively decomposing the system using product formulas to effectively reduce the norm. We obtain further improvements by simulating commuting Hamiltonian terms with an average combination cost. This gives simulations with complexities almost linear in the space-time volume for $\alpha \geq 2$ and improved scalings for $\alpha < 2$. We also implement faster Trotter steps when certain off-diagonal blocks of Hamiltonian coefficients exhibit low-rank properties. In particular, we achieve a nearly linear spacetime volume scaling for all $\alpha \geq 1$ using a recursive low-rank decomposition of the Hamiltonian.

We apply our methods to simulate electronic structure Hamiltonians in second quantization in real space. Combining with a tight Trotter error analysis, we show that $\left(\eta^{1/3}n^{1/3} + \frac{n^{2/3}}{\eta^{2/3}}\right)n^{1+o(1)}$ gates suffice to simulate uniform electron gas with n spin orbitals and η electrons, improving the best results from previous work. An analogous result holds when the external potential of nuclei is included from the Born-Oppenheimer approximation.

Performing faster Trotter steps for Hamiltonians with arbitrary coefficients is a challenging task in general. To confirm this intuition, we prove a gate-complexity lower bound. Specifically, we construct a class of 2-local Hamiltonians with commuting Pauli- Z terms whose coefficients take a continuum range of values. We show that these Hamiltonians require at least $\Omega(n^2)$ gates to simulate with accuracy $\epsilon = \Omega(1/\text{poly}(n))$ for time $t = \Omega(\epsilon)$; thus the best method one can hope for is to naively exponentiate all the $\Theta(n^2)$ terms. Our proof depends on a gate-efficient reduction from the approximate synthesis of diagonal unitaries within the Hamming weight-2 subspace, which we then address by adapting a volume-comparison technique from previous work. Our result thus suggests the use of structural properties of the target Hamiltonian as both necessary and sufficient to achieve lower gate complexity for implementing Trotter steps.

1.1 Quantum algorithms for quantum simulation

Simulating many-body physical systems is one of the most promising applications of digital quantum computers. Indeed, the idea of quantum computing as originally proposed by Feynman [39], Manin [88] and others is strongly motivated by quantum simulation. Efficient quantum simulations can be used to extract static and dynamical properties of physical systems, which has potential applications in various areas such as condensed matter physics [30], chemistry [24, 90, 93], and high-

energy physics [10]. Meanwhile, recent developments in quantum simulation algorithms have also provided technical tools that influenced the design of other quantum algorithms [3, 25, 44, 52, 76] and proofs of other results in areas beyond quantum computing [1, 5, 57, 73, 113].

There are many quantum algorithms one can use to perform quantum simulation. At a high level, these algorithms can be categorized according to their default input models. Common Hamiltonian input models include: (i) Linear Combinations of Hermitians (LCH), where the target Hamiltonian takes the form

$$H = \sum_{\gamma=1}^{\Gamma} H_{\gamma}, \quad (1)$$

with H_{γ} Hermitian and the exponentials $e^{-itH_{\gamma}}$ implementable on a quantum computer. Simulation algorithms that work in this input model include one based on product formulas [11, 79] (such as the Lie-Trotter formula and its higher-order extensions), as well as a more recent algorithm based on random sampling [23]; and (ii) Linear Combinations of Unitaries (LCU), where the target Hamiltonian takes the form

$$H = \sum_{\gamma=1}^{\Gamma} \beta_{\gamma} U_{\gamma}, \quad (2)$$

with U_{γ} unitary, $\beta_{\gamma} > 0$ and the controlled operators $|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes U_{\gamma}$ implementable on a quantum computer. There are also various algorithms working in the LCU model, including one based on implementing truncated Taylor series [13] and one based on qubitization [84]. For our purpose, we will mainly focus on the algorithm based on product formulas as well as the qubitization algorithm, which we review in more detail in [Section 2.2](#) and [Section 2.3](#).

Naturally, there is no silver bullet method that solves all simulation problems of interest with the optimal gate complexity. Choices of algorithms should thus be made on a case-by-case basis. There are in fact a few common desirable features shared among many algorithms mentioned above. For instance, it has been well known that many LCU approaches such as the Taylor-series algorithm and qubitization have complexities (nearly) linear in the simulation time and logarithmic in the inverse accuracy. But similar scalings can be achieved using product formulas as well by taking linear combinations of formulas with different step sizes and repetition numbers [86]. LCU-type algorithms also have the appealing feature that they can be used to design other functions of Hamiltonians [82], which applies to problems beyond quantum simulation [44, 89] such as solving linear systems of equations [52, 81], preparing ground states [40, 76, 100], and performing phase estimation [89, 102]. Again, these problems can be well-solved using algorithms in the LCH model (and can sometimes be more resource friendly) as demonstrated in recent work such as [37, 78].

However, there is one question concerning the gate complexities of implementing these approaches which has not been satisfactorily answered so far. View a sufficiently long-time Hamiltonian evolution as a concatenation of short-time steps, and consider the complexity of simulating each short evolution. Then a host of techniques are recently developed to implement LCU approaches with cost depending on the target system size as opposed to the number of Hamiltonian terms [75, 115]; a similar goal may be realized for the randomized method via an importance sampling [23, 28, 116]. By contrast, how to provably achieve a similar complexity for product formulas was somewhat under investigated: we review related work on the implementation of Trotter steps in [Section 1.2](#).

For an n -qubit system with 2-local interactions, one has that the total number of terms scales like $\Theta(n^2)$ but the system size is only linear in n . This comparison becomes $\Theta(n^{\kappa})$ versus n for a κ -local Hamiltonian. In light of this gap, it is natural to ask when and how one can implement product formulas with cost scaling better than the naive method. We address this question by presenting

necessary and sufficient conditions under which faster Trotter steps are possible. As an application, we give a simulation algorithm for electronic structure Hamiltonians in second quantization in real space with complexity $\left(\eta^{1/3}n^{1/3} + \frac{n^{2/3}}{\eta^{2/3}}\right)n^{1+o(1)}$ —asymptotically the fastest real-space simulation to date.

1.2 Previous related work

We now discuss prior work on implementing Trotter steps that are relevant to our paper.

First, there are previous studies on the so-called fast-forwardability of Hamiltonian evolution [6, 48, 70, 71]. Some of those techniques, such as introducing efficiently computable phases and diagonalizing quadratic Hamiltonians, can also be used to perform Trotter steps. However, their fundamental goal is quite different from ours. In quantum fast-forwarding, one is asked to simulate the target Hamiltonian for a sufficiently long time, and the goal is to reduce the scaling of time in the gate complexity (potentially at the cost of increasing the system-size scaling). In contrast, each Trotter step only approximates the ideal evolution for a short time, and so the scaling with time is no longer a key contribution to the complexity. Instead, our main goal here is to reduce the dependence on the size of the simulated system.

For electronic structure models represented under arbitrary basis, the number of terms in the Hamiltonian typically scales like $\Theta(n^4)$ for n spin orbitals, so a naive implementation of Trotter steps would have cost scaling $\Theta(n^4)$. To address this, recent work developed quantum circuits based on low-rank factorizations of such systems [94, 98] (see [4, 104] for more recent developments of such a method). Specifically, they apply product formulas to decompose the Hamiltonian into multiple components, each of which has coefficients with certain low-rank properties and can be further implemented by diagonalization. The gate complexity of the resulting circuits would then depend on the value of rank as opposed to the number of Hamiltonian terms, which significantly reduces the cost per time step. However, those work did not rigorously analyze the total complexity of the proposed methods, and it is unclear how much overall advantage their approach can offer. In fact, their factorization does not seem to preserve the commutation relations between Hamiltonian terms and could potentially introduce a Trotter error larger than the naive approach (so more Trotter steps would be required to reach the same simulation accuracy).

Another related approach to reducing the complexity of quantum simulation is to truncate Hamiltonian terms of small sizes. Such truncations are useful for not only performing Trotter steps [32, 33, 118], but also implementing more advanced quantum simulation algorithms [14, 58, 91]. Generally speaking, the error introduced in the truncation will grow linearly with time, so the simulation is accurate only when the evolution is sufficiently short. In particular, for rapidly decaying power-law interactions with exponent $\alpha > 2$, a truncation is possible only when $t = \mathcal{O}(n^{(\alpha-2)/(\alpha-1)})$ [32]. For simulations of chemistry and material models, truncation thresholds can often be determined empirically under certain assumptions of the model Hamiltonians.

Here, our work considers simulating 2-local Hamiltonians with interaction strength decaying according to power law, and we study the cost of implementing one short Trotter step as well as the entire long-time simulation, using the so-called block encoding technique and recursive/hierarchical low-rank decompositions [46, 50]. Our motivation for using the low-rank decomposition partly overlaps with that of a recent work by Nguyen, Kiani, and Lloyd [96], but the main problems we study are different. Instead of Hamiltonian simulation, they studied the block encoding of kernel matrices of the form

$$K = \sum_{j,k=1}^n \beta_{j,k} |j\rangle\langle k|, \quad (3)$$

where $\beta_{j,k}$ can be power-law functions such as $1/|j-k|^\alpha$. The matrix K is an n -dimensional operator and has spectral norm $\|K\| = \|\beta\| = \Theta(1)$ for $\alpha > 1$.¹ Instead, our problem centers around the simulation of

$$H = \sum_{j,k=1}^n \beta_{j,k} X_j Y_k, \quad (4)$$

which is a 2^n -dimensional operator and has spectral norm generally scaling with the *vector* 1-norm of coefficients: $\|H\| = \Theta(\|\beta\|_1) = \Theta(n)$ for $\alpha > 1$. Thus, a naive block encoding of our H will have an intrinsically worse normalization factor than that of their K ; see [Section 2.3](#) for further explanations of how such normalization factors affect the complexity of quantum simulation. Nevertheless, we overcome this technical obstacle by recursively decomposing the Hamiltonian using product formulas, which significantly reduces the 1-norm while maintaining the overall scaling of gate count.

Finally, we note that there is a large body of previous work analyzing and optimizing the concrete resources for implementing Trotter steps, for both near-term and fault-tolerant quantum computers (see Refs. [\[30, 64, 107\]](#) as well as other work citing and cited by these papers). We have not attempted to optimize the constant factors of the complexity of our methods, but we consider such optimizations to be an interesting subject for future investigation.

1.3 Faster Trotter steps

Consider a Hamiltonian H with Γ terms $H = \sum_{\gamma=1}^{\Gamma} H_{\gamma}$. If each Hamiltonian term can be exponentiated on a quantum computer with cost $\mathcal{O}(1)$, then one can simulate the evolution of H for a short time using product formulas, and the complexity would scale like $\mathcal{O}(\Gamma)$. We will identify scenarios in which improved implementations of Trotter steps are possible with gate complexities sublinear in Γ .

We focus on a class of 2-local Hamiltonians in one spatial dimension with magnitude of the coefficients decaying with distance x according to power law $1/x^\alpha$ ($\alpha > 0$). We describe how our results can be extended to higher spatial dimensions and potentially to more general systems in [Appendix C](#). We restrict to power-law models because product formulas are known to provide the fastest method for simulating this class of Hamiltonians, so we can directly compare our result with the state of the art. Examples of power-law interactions include the Coulomb interaction between charged particles and the dipole-dipole interaction between molecules, both of which are ubiquitous in quantum chemistry—a primary target application of quantum computation. In physics, impressive controls in recent experiments with trapped ions [\[19, 62\]](#), Rydberg atoms [\[105\]](#), and ultra cold atoms and polar molecules [\[38, 119\]](#) have enabled the possibility to study new phases of matter with power-law interactions [\[20, 35, 36, 45, 69, 92\]](#) and contributed to a growing interest in simulating such systems. In fact, we will describe a direct application of our method in [Section 6](#) for faster simulations of electronic structure Hamiltonians in real space.

Assume that the coefficients of the target Hamiltonian are efficiently computable. As explained above, there have been a host of techniques developed recently based on the notion of block encoding, which enables simulation with complexities depending only on the system size. One may ask if these techniques also lead to faster Trotter steps with a similar cost scaling. Unfortunately, the answer is negative in general. This is because a block encoding typically introduces a normalization

¹Roughly speaking, they are implementing operators in first quantization. One may attempt to use a similar approach on the Coulomb interaction to improve the electronic structure simulation. We are not aware of a simple realization of this idea. Existing study of this problem uses a computation model stronger than the circuit model [\[29\]](#).

factor proportional to the 1-norm of the Hamiltonian coefficients, and we thus need to repeat a corresponding number of times to perform Trotter steps. For instance, one can block-encode power-law Hamiltonians with gate complexity $\Theta(n)$, but this introduces a normalization factor proportional to the 1-norm which is generally $\Theta(n)$ for power-law interactions with $\alpha > 1$. Meanwhile, a Trotter step for the power-law models has an almost constant evolution time. So one roughly needs a total number of

$$\underbrace{n}_{\text{block encoding}} \cdot \underbrace{n}_{\text{effective time}} = \underbrace{n^2}_{\text{Trotter step complexity}} \quad (5)$$

gates to implement a single Trotter step, which has no benefit over the naive implementation. See [Section 3.1](#) for a more detailed explanation of this issue.²

We develop a method based on block encoding that overcomes the above technical issue. The key observation is that product formulas can be used to reduce the 1-norm of Hamiltonian coefficients “almost for free”: we apply product formulas to recursively decompose the Hamiltonian into multiple groups, but such a coarse-grained decomposition introduces a Trotter error no larger than the naive approach. We choose the decomposition to significantly reduce the 1-norm of each group while maintaining the overall scaling of the gate complexity, giving an efficient block encoding circuit. The resulting simulation has gate complexity $(nt)^{1+o(1)}$ when $\alpha \geq 2$, and $n^{3-\alpha+o(1)}t^{1+o(1)}$ when $\alpha < 2$. We formally state this theorem as [Theorem 1](#) in [Section 3](#) and preview it below.

Theorem (Faster Trotter steps using block encoding). *Consider 2-local Hamiltonians $H = \sum_{\sigma \neq \sigma' \in \{x,y,z\}} \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(\sigma,\sigma')} P_j^{(\sigma)} P_k^{(\sigma')}$, where $|\beta_{j,k}^{(\sigma,\sigma')}| \leq 1/|j-k|^\alpha$ for some constant $\alpha > 0$ and $P^{(\sigma)}$ ($\sigma = x, y, z$) are Pauli matrices. Let $t > 0$ be the simulation time and $\epsilon > 0$ be the target accuracy. Assume that the coefficient oracle*

$$O_\beta |\sigma, \sigma', j, k, 0\rangle = |\sigma, \sigma', j, k, \beta_{j,k}^{(\sigma,\sigma')}\rangle \quad (6)$$

can be implemented with gate complexity $\mathcal{O}(\text{polylog}(nt/\epsilon))$. Then H can be simulated with complexity

$$\begin{cases} \sqrt{\lambda_{\text{block}}} nt \left(\frac{nt}{\epsilon}\right)^{o(1)}, & \alpha \geq 2, \\ \sqrt{\lambda_{\text{block}}} n^{3-\alpha} t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & 0 < \alpha < 2. \end{cases} \quad (7)$$

Here, $\lambda_{\text{block}} \geq 1$ defined in Eq. (82) quantifies the deviation of the Hamiltonian coefficient distribution from a power law ($\lambda_{\text{block}} = \mathcal{O}(1)$ if the coefficient distribution exactly matches a power law).

The Hamiltonian decomposition we study has an additional feature that all the terms within each group commute with each other. Such terms can be simulated from block encodings of their subcomponents with an average combination cost, similar to the interaction-picture simulation [\[81, 87\]](#) but without an exponentially growing factor. We leverage this observation to further improve the complexity to $n^{2-\alpha/2+o(1)}t^{1+o(1)}$ for $1 \leq \alpha < 2$ and $n^{5/2-\alpha+o(1)}t^{1+o(1)}$ for $\alpha < 1$. We summarize our result in [Theorem 2](#) previewed in the following and describe the details in [Section 4](#).

Theorem (Faster Trotter steps using average-cost simulation). *Consider 2-local Hamiltonians $H = \sum_{\sigma \neq \sigma' \in \{x,y,z\}} \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(\sigma,\sigma')} P_j^{(\sigma)} P_k^{(\sigma')}$, where $|\beta_{j,k}^{(\sigma,\sigma')}| \leq 1/|j-k|^\alpha$ for some constant $\alpha > 0$*

²The fractional-query algorithm also implements a compressed version of Trotter steps using more advanced simulation techniques [\[12\]](#). That algorithm is superseded by the block-encoding method [\[72\]](#), so it suffers from the same normalization-factor issue pointed out here.

and $P^{(\sigma)}$ ($\sigma = x, y, z$) are Pauli matrices. Let $t > 0$ be the simulation time and $\epsilon > 0$ be the target accuracy. Assume that the coefficient oracle

$$O_\beta |\sigma, \sigma', j, k, 0\rangle = |\sigma, \sigma', j, k, \beta_{j,k}^{(\sigma, \sigma')}\rangle \quad (8)$$

can be implemented with gate complexity $\mathcal{O}(\text{polylog}(nt/\epsilon))$. Then H can be simulated with complexity

$$\begin{cases} \sqrt{\lambda_{\text{avg}}} n^{2-\alpha/2} t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & 1 \leq \alpha < 2, \\ \sqrt{\lambda_{\text{avg}}} n^{5/2-\alpha} t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & 0 < \alpha < 1. \end{cases} \quad (9)$$

Here, $\lambda_{\text{avg}} \geq 1$ defined in Eq. (109) quantifies the deviation of the Hamiltonian coefficient distribution from a power law ($\lambda_{\text{avg}} = \mathcal{O}(1)$ if the coefficient distribution exactly matches a power law).

We also realize faster Trotter steps when certain blocks of Hamiltonian coefficients exhibit low-rank properties. Such assumptions were studied in the context of block encoding kernel matrices [96]. Under the same hierarchical low-rank assumptions [46, 50], we directly implement the diagonalization procedure without using block encoding to achieve gate complexity $(nt)^{1+o(1)}$ nearly linear in the spacetime volume for all $\alpha \geq 1$. This is summarized in [Theorem 3](#) and restated below. See [Section 5](#) for details. We summarize our improvements (in a simplified form) in [Table 1](#) for simulating power-law Hamiltonians in general d spatial dimensions.

Theorem (Faster Trotter steps using low-rank decomposition). *Consider 2-local Hamiltonians $H = \sum_{\sigma \neq \sigma' \in \{x, y, z\}} \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(\sigma, \sigma')} P_j^{(\sigma)} P_k^{(\sigma')}$, where $|\beta_{j,k}^{(\sigma, \sigma')}| \leq 1/|j-k|^\alpha$ for some constant $\alpha > 0$ and $P^{(\sigma)}$ ($\sigma = x, y, z$) are Pauli matrices. Let $t > 0$ be the simulation time and $\epsilon > 0$ be the target accuracy. Then H can be simulated with complexity*

$$\begin{cases} \rho n t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & \alpha \geq 1, \\ \rho n^{2-\alpha} t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & 0 < \alpha < 1. \end{cases} \quad (10)$$

Here, ρ defined in Eq. (127) is the maximum rank of certain off-diagonal blocks of coefficient matrices ($\rho = \mathcal{O}(\log(nt/\epsilon))$ if the coefficient distribution exactly matches a power law in one spatial dimension).

Although we achieve various speedups for simulating power-law Hamiltonians based on different techniques, we have used recursion in the development of all our methods, and the core idea behind our improvements can all be understood through the so-called ‘‘master theorem’’ [34, 103]. Specifically, to solve a problem of size n using recursion, we divide the problem into m subproblems, each of which can be seen as an instance of the original problem of size n/m , so

$$\text{cost}_{\text{rec}}(n) = m \text{cost}_{\text{rec}}\left(\frac{n}{m}\right) + \text{cost}(n), \quad (11)$$

where $\text{cost}(n)$ quantifies the additional cost to combine solutions of the subproblems in the current layer of recursion and $\text{cost}_{\text{rec}}(n)$ denotes the total cost of the recursion. Then the master theorem asserts that, under certain assumptions of the cost function, the scaling of $\text{cost}_{\text{rec}}(n)$ is the same as that of $\text{cost}(n)$ up to a polylogarithmic factor. See [Lemma 1](#) of [Section 2.1](#) for a more formal description of this result. However, performing the combination step can often be much simpler than directly solving the full problem, and one then expects to get a better $\text{cost}_{\text{rec}}(n)$ when $\text{cost}(n)$ is improved. We show that improved recursions are indeed possible for power-law systems by reducing

Hamiltonian	$\alpha \geq 2d$	$d \leq \alpha < 2d$	$0 < \alpha < d$
Method			
Naive Trotter steps [32]	$n^{2+o(1)}t^{1+o(1)}$	$n^{2+o(1)}t^{1+o(1)}$	$n^{3-\frac{\alpha}{d}+o(1)}t^{1+o(1)}$
Block encoding ($\times\sqrt{\lambda_{\text{block}}}$)	$(nt)^{1+o(1)}$	$n^{3-\frac{\alpha}{d}+o(1)}t^{1+o(1)}$	$n^{3-\frac{\alpha}{d}+o(1)}t^{1+o(1)}$
Average-cost ($\times\sqrt{\lambda_{\text{avg}}}$)	—	$n^{2-\frac{\alpha}{2d}+o(1)}t^{1+o(1)}$	$n^{\frac{5}{2}-\frac{\alpha}{d}+o(1)}t^{1+o(1)}$
Low-rank ($\times\rho$)	$(nt)^{1+o(1)}$	$(nt)^{1+o(1)}$	$n^{2-\frac{\alpha}{d}+o(1)}t^{1+o(1)}$

Table 1: Comparison of our results and the best previous results for simulating power-law interactions $1/x^\alpha$ in d spatial dimensions. Upper bounds are used for some gate complexity expressions for presentational purpose. The truncation result of [32] only holds for a sufficiently short time and is thus not compared here. Both the block encoding and the average-cost simulation method assume an oracular access to the coefficients of the Hamiltonian. They also introduce factors $\sqrt{\lambda_{\text{block}}}$ and $\sqrt{\lambda_{\text{avg}}}$ that are constant if Hamiltonian coefficients have an exact power-law distribution. The low-rank method has a dependence on ρ which is polylogarithmic in the input parameters for all power-law models in 1D and many power-law models such as Coulomb interactions in higher spatial dimensions.

the normalization factor of the Hamiltonian and by exploiting low-rank properties of certain blocks of the Hamiltonian coefficients, leading to faster quantum simulation by recursion.

The electronic structure Hamiltonian is one of the most widely studied candidate models in quantum simulation [24, 90, 93]. An efficient simulation of such Hamiltonians could provide insights to various problems in chemistry and material science. Here, we focus on a simulation in real space, an idea investigated by Kassal et al. [61] and subsequently pursued by later work such as [9, 26, 60, 66, 110]. Although the full Hamiltonian does not satisfy power laws, the Coulomb potential part can be represented in second quantization with the magnitude of coefficients decaying as $1/x$, to which our method applies. In particular, we can choose $\rho = \mathcal{O}(\log(nt/\epsilon))$ in the low-rank decomposition [47] to efficiently implement Trotter steps. Combining with an improved Trotter error analysis, we show that $\left(\eta^{1/3}n^{1/3} + \frac{n^{2/3}}{\eta^{2/3}}\right)n^{1+o(1)}$ gates suffice to simulate the uniform electron gas with n spin orbitals and η electrons; we obtain an analogous result when the external potential of nuclei is introduced under the Born-Oppenheimer approximation. This improves the best previous results for simulating electronic structures. We describe this application in Section 6, with the improved Trotter error analysis detailed in Appendix B.

1.4 Circuit lower bound

It is worth noting that all our above methods hold under certain additional assumptions on the target Hamiltonian: we have assumed that either the Hamiltonian coefficients are efficiently computable or certain blocks of them have low rank. If no such structural properties are available, we are then faced with Hamiltonians with arbitrary coefficients, and intuitively there would be no implementation of Trotter steps better than the naive method.

We prove a circuit lower bound to justify this intuition. Specifically, we consider a class of 2-local Hamiltonians of the form

$$H = \sum_{1 \leq j < k \leq n} \beta_{j,k} Z_j Z_k, \quad (12)$$

where Z_j denotes the Pauli- Z operator on the j th qubit and coefficients are arbitrarily chosen from a continuum range of values $|\beta_{j,k}| \leq t$. Subclasses of such Hamiltonians are of interest in areas beyond quantum simulation [51]. Even for such commuting H , we show the gate-complexity lower bound $\Omega(n^2/\log(b|\mathcal{K}|))$ to evolve with accuracy $\epsilon = \Omega(1/\text{poly}(n))$ for $t = \Omega(\epsilon)$, using quantum

circuits of $b \geq n$ qubits with a gate set \mathcal{K} of finite size $|\mathcal{K}|$. For circuits with a continuous gate set, we can first compile them using a finite universal gate set (say applying the Clifford+T synthesis [17] or the Solovay-Kitaev theorem [63]) and invoke the above bound. We discuss the circuit lower bounds in detail in Section 7 with the result summarized in Corollary 1 and previewed here.

Corollary (Simulating 2-local commuting Hamiltonians). *Consider 2-local Hamiltonians $H = \sum_{1 \leq j < k \leq n} \beta_{j,k} Z_j Z_k$, where coefficients take values up to $|\beta_{j,k}| \leq t$. Given accuracy $0 < \epsilon < 1/3$, number of qubits $b \geq n$, and 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$, if $t = \Omega(\epsilon)$,*

$$\begin{aligned} & \min \{g \mid \forall \text{ 2-local Hamiltonian } H, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from } \mathcal{K}, \|e^{-iH} - V\| \leq \epsilon\} \\ & = \Omega \left(\frac{n^2}{\log(b|\mathcal{K}|)} - n \text{polylog} \left(\frac{n}{\epsilon} \right) \right). \end{aligned} \tag{13}$$

Under the same assumption but choosing \mathcal{K} to be the set of arbitrary 2-qubit gates,

$$\begin{aligned} & \min \{g \mid \forall \text{ 2-local Hamiltonian } H, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from } \mathcal{K}, \|e^{-iH} - V\| \leq \epsilon\} \\ & = \Omega \left(\frac{n^2}{\log b} - n \text{polylog} \left(\frac{n}{\epsilon} \right) \right). \end{aligned} \tag{14}$$

Underpinning our circuit lower bound proof is an efficient reduction from the approximate synthesis of diagonal unitaries within the Hamming weight-2 subspace, up to a gate overhead of $\mathcal{O}(n \text{polylog}(n/\epsilon))$. We describe this reduction in Section 7.2 (with the circuit illustrated in Figure 4). The Hamming weight-2 subspace has dimensions $\binom{n}{2} = \Theta(n^2)$, so our problem is reduced to studying diagonal operators

$$\mathcal{D}_{\theta_{\max}} = \left\{ \sum_x e^{i\theta_x} |x\rangle\langle x|, |\theta_x| \leq \theta_{\max} \right\} \tag{15}$$

for $\Theta(n^2)$ values of x . We then show in Section 7.1 that such diagonal unitaries require roughly n^2 gates to approximately implement, generalizing a previous lower bound for exact synthesis due to Bullock and Markov [21, 117]. See also recent work [59] for a related bound expressed in terms of measure. Our argument is based on an adaption of a technique of Knill [67]. Knill proved asymptotic circuit lower bounds for synthesizing the full unitary group using a volume-comparison technique, but here we only consider diagonal unitaries whose volume can be easily evaluated in closed form (see Section 2.4). The volume of diagonal unitaries predominantly depends on the dimensionality $\Theta(n^2)$, which leads to our desired lower bound scaling.

Additionally, we prove a second lower bound in Appendix A showing that any approximate realization of the coefficient oracle O_β requires $\Omega(n^2)$ gates in the worst case. Combining both our upper and lower bounds, we conclude that the use of structural properties of the Hamiltonian is both necessary and sufficient to achieve lower gate complexity for implementing Trotter steps.

We briefly summarize in Section 2 the preliminaries of our paper and present in Section 8 a collection of questions related to our result for future work.

2 Preliminaries

2.1 Notation and terminology

We now introduce notation and terminology to be used in the remainder of our paper.

We use lowercase Latin letters as well as the Greek alphabet (in both upper and lower cases) to denote scalars, and we save uppercase Latin letters for matrices and operators. For instance, we will use n for the system size, t for the evolution time (assuming $t > 0$ without loss of generality), ϵ for the simulation accuracy, and we write H to denote the target Hamiltonian, X, Y, Z to denote Pauli operators, and I to denote the identity matrix. When discussing fermionic Hamiltonians, we use A^\dagger, A , and N to represent fermionic creation, annihilation, and occupation-number operators respectively [111]. We write the commutator of matrices B and C as $[B, C] := BC - CB$ when the multiplications are well-defined. When multiplying noncommutative quantities, we use abbreviations like $\prod_{j=1}^n$ to denote the ordering where the smallest index appears on the right, e.g., $\prod_{j=1}^n U_j = U_n \cdots U_1$. We let a summation be zero and a product be one if their lower limits exceed the upper limits. We interchangeably use the decimal representation $|x\rangle$ ($x = 0, \dots, 2^n - 1$) and the binary representation $|x_{n-1}, \dots, x_0\rangle$ ($x_j = 0, 1$) of computational basis states if no ambiguity arises.

We also construct vectors of scalars and operators and use subscripts to index them, e.g. $\beta_{j,k}$ and Z_j . We then define the transpose operation $\beta_{j,k}^\top := \beta_{k,j}$. Our focus will be on simulating 2-local systems throughout the paper, so the Hamiltonian coefficients will typically have no more than two subscript indices. We assume that coefficients in an n -qubit Hamiltonian can be represented in binary using $\mathcal{O}(\log(nt/\epsilon))$ bits, for otherwise we may truncate the binary representation and simulate the truncated Hamiltonian with error at most $\mathcal{O}(\epsilon)$.

Our analysis requires various norms defined for vectors and matrices. For $[\beta_{j,k}]$, we define the vector 1-norm $\|\beta\|_1 := \sum_{j,k} |\beta_{j,k}|$, the max-norm $\|\beta\|_{\max} := \max_{j,k} |\beta_{j,k}|$, the Euclidean norm $\|\beta\| := \sqrt{\sum_{j,k} |\beta_{j,k}|^2}$ and the induced 1-norm $\|\|\beta\|\|_1 := \max_j \sum_k |\beta_{j,k}|$. We also need a restricted version of the induced 1-norm defined as

$$\|\|\beta\|\|_{1, [\eta]} := \max_j \max_{k_1 < \dots < k_\eta} (|\beta_{j,k_1}| + \dots + |\beta_{j,k_\eta}|). \quad (16)$$

By definition, this norm $\|\|\beta\|\|_{1, [\eta]}$ only sums the largest η elements in a row (maximized over all rows), and is thus always upper bounded by the induced 1-norm. We will see later in the quantum chemistry application that the gap between these two norms can be significant. We use $\|B\|$ to denote the operator norm of B ; this is also known as the spectral norm and its value is given by the largest singular value of B .

We will use calligraphic uppercase letters to denote (un)structured sets. For instance, we use \mathcal{H} to represent an arbitrary finite-dimensional Hilbert space with (normalized) quantum states $|\psi\rangle \in \mathcal{H}$ (we will write \mathbb{C}^m or \mathbb{R}^m if the dimensionality m is explicitly provided). Given an underlying n -qubit system, we denote \mathcal{W}_η to be the subspace spanned by computational basis states with Hamming weight η ($\dim(\mathcal{W}_\eta) = \binom{n}{\eta}$); \mathcal{W}_η also denotes the subspace spanned by states with η particles for second-quantized fermionic systems. We may then define the operator norm restricted to the Hamming weight- η subspace

$$\|B\|_{\mathcal{W}_\eta} := \max_{|\phi_\eta\rangle, |\psi_\eta\rangle \in \mathcal{W}_\eta} |\langle \phi_\eta | B | \psi_\eta \rangle| \quad (17)$$

for an arbitrary n -qubit operator B . Given $[\beta_{j,k}]$ and a collection $\mathcal{B} \subseteq \mathbb{Z} \times \mathbb{Z}$ of pairs of indices (u, v) , we define the restricted max-norm and 1-norm as

$$\|\beta\|_{\max, \mathcal{B}} := \max_{(u,v) \in \mathcal{B}} |\beta_{u,v}|, \quad \|\beta\|_{1, \mathcal{B}} := \sum_{(u,v) \in \mathcal{B}} |\beta_{u,v}|. \quad (18)$$

And as introduced earlier, the set $\mathcal{D}_{\theta_{\max}}$ is the set of diagonal unitaries with phase angles between $-\theta_{\max}$ and θ_{\max} .

We say an operator $G : \mathcal{G} \rightarrow \mathcal{H}$ is an *isometry* if $G^\dagger G = I$. By definition, G is necessarily injective and G^\dagger is necessarily surjective, whereas GG^\dagger is an orthogonal projection on \mathcal{H} with image $\text{im}(GG^\dagger) = \text{im}(G)$ and kernel $\ker(GG^\dagger) = \ker(G^\dagger)$. We thus obtain the Hilbert space isomorphism $\mathcal{G} \cong \text{im}(GG^\dagger) \subseteq \mathcal{H}$ specified by the operators G and G^\dagger . Choosing any orthonormal basis that respects this isomorphism, we have the matrix representation

$$G = \begin{bmatrix} I \\ 0 \end{bmatrix}, \quad G^\dagger = [I \quad 0]. \quad (19)$$

Examples of isometries include: (i) unitary operators U ; (ii) quantum states $|\psi\rangle$; (iii) tensor product $G_1 \otimes G_2$ if G_1 and G_2 are isometries; and (iv) composition $G_2 G_1$ if G_1 and G_2 are isometries and the composition is well defined. Isometries will be used later in [Section 2.3](#) to describe block encodings and the qubitization algorithm.

Finally, we use $\mathcal{O}(\cdot)$ and $\Omega(\cdot)$ to mean asymptotically bounded above and below respectively, write $\Theta(\cdot)$ if both relations hold, and use the tilde symbol to suppress polylogarithmic factors. This is similar to the notation of a previous work on Trotter error analysis [\[32\]](#), except we do not need their $O(\cdot)$ for order conditions. To analyze the scaling of functions satisfying recurrence relations, we use the following adapted version of the master theorem [\[34, 103\]](#).

Lemma 1 (Master theorem). *Let $\text{cost}, \text{cost}_{\text{rec}} : \mathbb{Z}_{\geq 1} \rightarrow \mathbb{R}_{\geq 0}$ be nonnegative functions defined for positive integers, such that there exist constant $\gamma \geq 0$, $m \in \mathbb{Z}_{\geq 2}$, and $n_0 \in \mathbb{Z}_{\geq 1}$ for which*

$$\begin{cases} \text{cost}_{\text{rec}}(n) \leq \gamma, & 1 \leq n \leq n_0, \\ \text{cost}_{\text{rec}}(n) \leq m \text{cost}_{\text{rec}}(\lceil \frac{n}{m} \rceil) + \text{cost}(n), & n > n_0. \end{cases} \quad (20)$$

If $\text{cost}(n) = \mathcal{O}(n^c \log^k(n))$ for some $c \geq 0$ and $k \in \mathbb{Z}_{\geq 0}$, then

$$\text{cost}_{\text{rec}}(n) = \begin{cases} \mathcal{O}(n), & 0 \leq c < 1, \\ \mathcal{O}(n \log^{k+1}(n)), & c = 1, \\ \mathcal{O}(n^c \log^k(n)), & c > 1. \end{cases} \quad (21)$$

2.2 Product formulas

Consider a Hamiltonian given in the LCH form $H = \sum_{\gamma=1}^{\Gamma} H_\gamma$. We can well-approximate the evolution under H for a short time using *product formulas* with error high-order in time. A longer evolution can then be simulated by repeating the short-time steps. Product formulas provide a simple yet surprisingly efficient approach to quantum simulation. Indeed, recent work has shown that product formulas: (i) can simulate geometrically local lattice systems [\[49\]](#) with nearly optimal gate complexity [\[31\]](#); (ii) have the lowest asymptotic cost for electronic structure Hamiltonians in second quantization in the plane wave basis [\[111\]](#); and (iii) are advantageous for simulating general κ -local Hamiltonians [\[32\]](#) (although our work achieves further improvements regarding this last point). In addition, product formulas have also been widely used in classical simulations of quantum systems and in areas beyond quantum computing [\[16\]](#).

By definition, the cost of the product-formula approach is determined by both the repetition number of short-time steps and the complexity of implementing each step. To elaborate, first consider a simple example where we use the Lie-Trotter formula $e^{-itB}e^{-itA}$ to simulate a 2-term Hamiltonian $H = A + B$. Then it holds that

$$e^{-itB}e^{-itA} - e^{-it(A+B)} = \int_0^t d\tau_1 \int_0^{\tau_1} d\tau_2 e^{-i(t-\tau_1)(A+B)} e^{-i\tau_1 B} e^{i\tau_2 B} [iB, iA] e^{-i\tau_2 B} e^{-i\tau_1 A}, \quad (22)$$

which implies the *Trotter error* bound

$$\left\| e^{-itB} e^{-itA} - e^{-it(A+B)} \right\| \leq \frac{t^2}{2} \|[A, B]\|. \quad (23)$$

This approximation error is small for a sufficiently short-time *Trotter step*. For a longer simulation, we apply r repetitions of the same step with time t/r , obtaining

$$\left\| \left(e^{-i\frac{t}{r}B} e^{-i\frac{t}{r}A} \right)^r - e^{-it(A+B)} \right\| \leq r \left\| e^{-i\frac{t}{r}B} e^{-i\frac{t}{r}A} - e^{-i\frac{t}{r}(A+B)} \right\| \leq \frac{t^2}{2r} \|[A, B]\|. \quad (24)$$

To achieve an error at most ϵ , it thus suffices to take

$$r = \left\lceil \frac{\|[A, B]\| t^2}{2\epsilon} \right\rceil. \quad (25)$$

We now have a total of $2r$ elementary exponentials, where r is the number of Trotter steps determined by the error bound Eq. (23).

In general, we can simulate the Hamiltonian $H = \sum_{\gamma=1}^{\Gamma} H_{\gamma}$ using a p th-order product formula $S_p(t)$, where p is a positive integer that can be arbitrarily large. Similar to above, we need to bound the cost of implementing Trotter steps, as well as the repetition number r which is in turn determined by the analysis of Trotter error. We introduce the following Trotter error bound with a commutator scaling established by [32]. Although Trotter error analysis can be further tightened using additional assumptions of the quantum simulation problem [2, 22, 27, 53, 55, 68, 74, 106, 121], those improvements are not relevant to our results and will not be further discussed here.

Lemma 2 (Trotter error with commutator scaling). *Let $H = \sum_{\gamma=1}^{\Gamma} H_{\gamma}$ be a Hamiltonian in the LCH form and $S_p(t)$ be a p th-order product formula with respect to this decomposition. We have*

$$\|S_p(t) - e^{-itH}\| = \mathcal{O} \left(\sum_{\gamma_1, \gamma_2, \dots, \gamma_{p+1}=1}^{\Gamma} \|[H_{\gamma_{p+1}}, \dots, [H_{\gamma_2}, H_{\gamma_1}]]\| t^{p+1} \right). \quad (26)$$

Furthermore, if H is a fermionic Hamiltonian in second quantization and H_{γ} are number preserving,

$$\|S_p(t) - e^{-itH}\|_{\mathcal{W}_{\eta}} = \mathcal{O} \left(\sum_{\gamma_1, \gamma_2, \dots, \gamma_{p+1}=1}^{\Gamma} \|[H_{\gamma_{p+1}}, \dots, [H_{\gamma_2}, H_{\gamma_1}]]\|_{\mathcal{W}_{\eta}} t^{p+1} \right). \quad (27)$$

We denote the sum of nested commutators as

$$\tilde{\alpha}_{\text{comm}} := \sum_{\gamma_1, \gamma_2, \dots, \gamma_{p+1}=1}^{\Gamma} \|[H_{\gamma_{p+1}}, \dots, [H_{\gamma_2}, H_{\gamma_1}]]\|. \quad (28)$$

Then to simulate for time t with accuracy ϵ , it suffices to take $r = \mathcal{O} \left(\tilde{\alpha}_{\text{comm}}^{1/p} t^{1+1/p} / \epsilon^{1/p} \right)$, which simplifies to

$$r = \frac{\tilde{\alpha}_{\text{comm}}^{o(1)} t^{1+o(1)}}{\epsilon^{o(1)}} \quad (29)$$

by choosing p sufficiently large. For the fermionic case, we can simply replace the spectral norm $\|\cdot\|$ by its restriction to the Hamming weight- η subspace $\|\cdot\|_{\mathcal{W}_{\eta}}$.

A naive implementation of Trotter steps has an asymptotic cost of $\mathcal{O}(\Gamma)$, giving total gate complexity $\mathcal{O}(r\Gamma)$. The purpose of this work is to identify conditions under which the cost of Trotter steps does not explicitly scale with Γ . We will present three methods, one based on the block encoding technique (Section 3), one based on an average-cost simulation (Section 4) and one based on a recursive low-rank decomposition (Section 5). We review the basic notion of block encoding as well as the related qubitization algorithm in the next subsection.

2.3 Block encoding and qubitization

Block encodings, together with the related qubitization algorithm, provide a versatile framework for simulating Hamiltonians in the LCU form and beyond. Such techniques enable quantum simulation with a complexity linear in the evolution time and logarithmic in the inverse accuracy [83], and can be extended to solve problems other than simulation [44, 89] (although these goals are sometimes also achievable via product formulas as explained in Section 1.1). In addition, these techniques have been utilized by recent work to design classical algorithms for simulating quantum systems [41].

To explain the idea of block encoding and qubitization, we will use the notion of isometries. Specifically, consider two isometries $G_0 : \mathcal{G}_0 \rightarrow \mathcal{H}$, $G_1 : \mathcal{G}_1 \rightarrow \mathcal{H}$ and a unitary $U : \mathcal{H} \rightarrow \mathcal{H}$. We say an operator $B : \mathcal{G}_0 \rightarrow \mathcal{G}_1$ is *block-encoded* by G_0 , G_1 , and U if

$$B = G_1^\dagger U G_0. \quad (30)$$

Of course, this unitary dilation is mathematically valid if and only if $\|B\| \leq 1$ [56, 2.7.P2]. However, there are many scenarios where additional normalization factors will be introduced when such block encodings are realized by quantum circuits. If we choose two bases with respect to the orthogonal decompositions $\mathcal{H} = \text{im}(G_0 G_0^\dagger) \oplus \text{im}(I - G_0 G_0^\dagger) = \text{im}(G_1 G_1^\dagger) \oplus \text{im}(I - G_1 G_1^\dagger)$, then U has a matrix representation that looks like

$$U = \begin{bmatrix} B & \cdot \\ \cdot & \cdot \end{bmatrix}, \quad (31)$$

where B is exhibited as the top-left matrix block; hence the name *block encoding*.

For quantum simulation, we fix $\mathcal{G}_0 = \mathcal{G}_1 = \mathcal{G}$ to be the target system space and we let $G_1^\dagger U G_0$ be Hermitian to represent physical Hamiltonians. Then, we have the following qubitization algorithm to perform quantum simulation of block-encoded operators [84].

Lemma 3 (Hamiltonian simulation by qubitization). *Let $G_0, G_1 : \mathcal{G} \rightarrow \mathcal{H}$ be isometries and $U : \mathcal{H} \rightarrow \mathcal{H}$ be a unitary such that $G_1^\dagger U G_0$ is Hermitian. Given a target evolution time t and accuracy ϵ , there exists a unitary $V_\varphi : \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathcal{H} \rightarrow \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathcal{H}$ parameterized by angles $\varphi_1, \dots, \varphi_r$ such that*

$$\left\| \left(\langle + | \otimes \frac{\langle 0 | \otimes G_0^\dagger + \langle 1 | \otimes G_1^\dagger}{\sqrt{2}} \right) V_\varphi \left(| + \rangle \otimes \frac{| 0 \rangle \otimes G_0 + | 1 \rangle \otimes G_1}{\sqrt{2}} \right) - e^{-it G_1^\dagger U G_0} \right\| \leq \epsilon. \quad (32)$$

The number of steps r is an even integer with the asymptotic scaling

$$r = \mathcal{O} \left(t + \log \left(\frac{1}{\epsilon} \right) \right), \quad (33)$$

and V_φ is explicitly defined as

$$\begin{aligned} V_\varphi &:= \prod_{k=0}^{r/2-1} \left(V_{\varphi_{2k+2}}^\dagger V_{\varphi_{2k+1}} \right) \\ V_{\varphi_j} &:= \left(e^{-i\varphi_j Z/2} \otimes I \right) \left(|+\rangle\langle +| \otimes I + |-\rangle\langle -| \otimes V \right) \left(e^{i\varphi_j Z/2} \otimes I \right) \\ V &:= (X \otimes I) \left(|0\rangle\langle 0| \otimes U + |1\rangle\langle 1| \otimes U^\dagger \right) \left(I - 2 \frac{|0\rangle \otimes G_0 + |1\rangle \otimes G_1}{\sqrt{2}} \frac{\langle 0| \otimes G_0^\dagger + \langle 1| \otimes G_1^\dagger}{\sqrt{2}} \right). \end{aligned} \quad (34)$$

Note that by our definition of \mathcal{O} , the complexity expression Eq. (33) is valid for t sufficiently large and ϵ sufficiently small. This expression may be modified if other parameter regimes are of interest.

To simulate an n -qubit Hamiltonian $H = \sum_{\gamma=1}^{\Gamma} \beta_{\gamma} U_{\gamma}$ in the LCU form, we define

$$\begin{aligned} G : \mathbb{C}^{2^n} &\rightarrow \mathbb{C}^{\Gamma} \otimes \mathbb{C}^{2^n}, & G &= \frac{1}{\sqrt{\|\beta\|_1}} \sum_{\gamma=1}^{\Gamma} \sqrt{\beta_{\gamma}} |\gamma\rangle \otimes I, \\ U : \mathbb{C}^{\Gamma} \otimes \mathbb{C}^{2^n} &\rightarrow \mathbb{C}^{\Gamma} \otimes \mathbb{C}^{2^n}, & U &= \sum_{\gamma=1}^{\Gamma} |\gamma\rangle\langle\gamma| \otimes U_{\gamma}, \end{aligned} \tag{35}$$

which implies

$$G^{\dagger} U G : \mathbb{C}^{2^n} \rightarrow \mathbb{C}^{2^n}, \quad G^{\dagger} U G = \frac{H}{\|\beta\|_1}. \tag{36}$$

We thus see that the target Hamiltonian is block-encoded by G (a state *preparation subroutine*) and U (an operator *selection subroutine*), but with a normalization factor that depends on the vector 1-norm $\|\beta\|_1$, which slows down quantum simulation. To approximate e^{-itH} , we then need to invoke Lemma 3 with an effective evolution time of $\|\beta\|_1 t$. This implies that the number of steps will now scale as

$$r = \mathcal{O} \left(\|\beta\|_1 t + \log \left(\frac{1}{\epsilon} \right) \right). \tag{37}$$

The preparation of a Γ -dimensional state naively requires $\mathcal{O}(\Gamma)$ gates to implement [108] and the selection of Γ terms has cost $\mathcal{O}(\Gamma)$ [30, Appendix G.4], assuming each controlled- U_{γ} has constant cost. This gives a total complexity of $\mathcal{O}(r\Gamma)$ for qubitization.

Having introduced the qubitization algorithm, we now make a few remarks about its circuit implementation.

First, qubitization approximately simulates the time evolution with a certain probability, and both the approximation error and the success probability can be analyzed using the condition Eq. (32). Indeed, an argument based on the triangle inequality shows that the qubitization algorithm succeeds with probability at least $(1 - \epsilon)^2 \geq 1 - 2\epsilon$, and the post-measurement operator has error at most $\sqrt{1 + \epsilon} - \sqrt{1 - \epsilon} \leq \sqrt{2}\epsilon$ [30, Lemma G.4 and Appendix H.1]. It therefore suffices to adjust the value of ϵ to meet the desired error and probability requirements.

Second, in our above LCU example, we have defined G and U that faithfully block-encode $H / \|\beta\|_1$. More realistically, we may consider state preparation and operator selection circuits that have lower cost but introduce error to the block encoding. Specifically, suppose $\tilde{G}^{\dagger} \tilde{U} \tilde{G} = \tilde{H} / \|\tilde{\beta}\|_1$, then we are effectively block encoding an erroneous Hamiltonian with error growing at most linearly in time: $\left\| e^{-it\|\tilde{\beta}\|_1 \tilde{G}^{\dagger} \tilde{U} \tilde{G}} - e^{-itH} \right\| \leq t \left\| H - \tilde{H} \right\|$. And we can suppress this error accordingly by increasing the accuracy of block encoding.

We will focus on the complexity of the state preparation and operator selection subroutines for the remainder of our paper. There is also an additional cost in qubitization to implement the operations in between these subroutines, but that cost is typically logarithmic in Γ and only makes a mild contribution to the total gate complexity. In particular, one can check that this holds for our block encodings of power-law Hamiltonians to be described in Section 3 and Section 4. So there is no loss of rigor to ignore this subdominant contribution in our analysis.

Finally, we point out that a host of techniques have been recently developed to improve the circuit implementation of block encoding and qubitization over the naive approach. For instance, if the target Hamiltonian consists of tensor product of Pauli operators, then the selection subroutine

$U = U^\dagger$ is both unitary and Hermitian, so we can reduce 1 ancilla from the qubitization circuit in [Lemma 3](#). Even when this Hermitian condition is not satisfied, one may still improve the implementation of the bidirectional control $|0\rangle\langle 0| \otimes U + |1\rangle\langle 1| \otimes U^\dagger$ using explicit structures of U . Also, the qubitization algorithm we have introduced in [Lemma 3](#) is mainly for dynamical simulation. When qubitization is used as a subroutine in quantum phase estimation, we only need to implement the operator V in Eq. (34) as opposed to the full circuit [15, 99].³ There are also some flexibilities about block encoding the LCU Hamiltonian beyond the standard approach of Eq. (35), such as allowing garbage registers [85], performing unary encodings [13], or using even more advanced circuit compilation tricks [109]. Additionally, given a block encoding $G_0, G_1 : \mathcal{G} \rightarrow \mathcal{H}$ and $U : \mathcal{H} \rightarrow \mathcal{H}$, there always exists unitary $W : \mathcal{H} \rightarrow \mathcal{H}$ such that $WG_1 = G_0$. Any such W satisfies $G_1^\dagger U G_0 = G_0^\dagger (WU) G_0$ and hence gives an alternative block encoding with only one isometry, simplifying the circuit implementation.

But perhaps the most significant improvement comes from the following observation: terms in the LCU Hamiltonian are sometimes indexed by vectors with a product structure, i.e.,

$$U_\gamma = U_{\gamma_1} \cdots U_{\gamma_s}, \quad \gamma_j = 1, \dots, \Gamma^{1/s}. \quad (38)$$

In this case, the selection subroutine takes a product form

$$\sum_{\gamma=1}^{\Gamma} |\gamma\rangle\langle\gamma| \otimes U_\gamma = \sum_{\gamma_1, \dots, \gamma_s=1}^{\Gamma^{1/s}} |\gamma_1 \cdots \gamma_s\rangle\langle\gamma_1 \cdots \gamma_s| \otimes U_{\gamma_1} \cdots U_{\gamma_s} \quad (39)$$

and can be implemented with complexity $\mathcal{O}(s\Gamma^{1/s})$ as opposed to $\mathcal{O}(\Gamma)$. Meanwhile, we can sometimes exploit the structure of Hamiltonian coefficients to also perform the preparation subroutine faster than the naive approach. This then yields a step of qubitization with complexity sublinearly in the number of terms Γ .

However, the above discussion ignored the scaling of the normalization factor $\|\beta\|_1$, which determines the number of repetition steps and is still the bottleneck of qubitization. In fact, the reduction of the 1-norm has been the central technical problem studied by many recent work of quantum simulation [75, 80, 115, 120]. We show in [Section 3](#) that this 1-norm can be significantly reduced using product formulas “almost for free”, which leads to new methods for simulating power-law Hamiltonians that improve the state of the art.

2.4 Volume of diagonal unitaries

In this section, we review techniques for computing the volume of a parameterized object in a high-dimensional Euclidean space. In particular, we compute the volume of diagonal unitaries with phase angles taking a continuum range of values between $-\theta_{\max}$ and θ_{\max} , which will be used later in [Section 7](#) to prove our circuit lower bound. These calculations can be seen as generalizations of line and surface integrals defined for low-dimensional spaces.

Let $\mathcal{M} \subset \mathbb{R}^m$ be a subset and $f : \mathcal{M} \rightarrow \mathbb{R}^q$ ($m \leq q$) be a mapping. Here, $f = [f_1 \cdots f_q]$ consists of q scalar functions and we can think of it as a parameterized description of a high-dimensional geometrical object, where \mathcal{M} is the set of all possible parameters. Then, given a subset of parameters $\mathcal{A} \subset \mathcal{M}$, we may compute the volume of $f(\mathcal{A})$ as

$$\text{vol}(f(\mathcal{A})) = \int_{\mathcal{A}} d\theta \sqrt{\det(G_f)}, \quad (40)$$

³This simplification follows from a spectral property of V : there exists an orthonormal basis under which V is exhibited as a direct sum of 1-by-1 and 2-by-2 matrices, whose eigenvalues correspond to eigenvalues of the target Hamiltonian (which is enough for phase estimation). We can then visualize V as a collection of single-qubit unitaries acting on orthogonal subspaces; hence the name “qubitization”.

where G_f is the Gramian whose element in the x th row and y th column is given by the inner product

$$\begin{bmatrix} \frac{\partial f_1}{\partial \theta_x} & \cdots & \frac{\partial f_q}{\partial \theta_x} \end{bmatrix} \begin{bmatrix} \frac{\partial f_1}{\partial \theta_y} \\ \vdots \\ \frac{\partial f_q}{\partial \theta_y} \end{bmatrix}. \quad (41)$$

For the above formula to hold mathematically, we require that \mathcal{M} is open, \mathcal{A} is Lebesgue measurable, and f is continuously differentiable and injective on \mathcal{A} . We refer the reader to standard analysis textbooks for a detailed explanation of this formula [42, 43].

Now, consider the set of diagonal unitaries

$$\mathcal{D}_{\theta_{\max}} := \left\{ \sum_{x=0}^{m-1} e^{i\theta_x} |x\rangle\langle x|, |\theta_x| \leq \theta_{\max} < \pi \right\}. \quad (42)$$

We may view them as an object in the Euclidean space \mathbb{R}^{2m} by choosing the parameterization

$$f(\theta_0, \dots, \theta_{m-1}) = [\cos \theta_0 \quad \sin \theta_0 \quad \cdots \quad \cos \theta_{m-1} \quad \sin \theta_{m-1}]. \quad (43)$$

Then, the Gramian matrix is simply the identity matrix

$$G_f = \begin{bmatrix} -\sin \theta_0 & \cos \theta_0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & -\sin \theta_1 & \cos \theta_1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & -\sin \theta_{m-1} & \cos \theta_{m-1} \end{bmatrix} \begin{bmatrix} -\sin \theta_0 & 0 & \cdots & 0 \\ \cos \theta_0 & 0 & \cdots & 0 \\ 0 & -\sin \theta_1 & \cdots & 0 \\ 0 & \cos \theta_1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & -\sin \theta_{m-1} \\ 0 & 0 & \cdots & \cos \theta_{m-1} \end{bmatrix} = I, \quad (44)$$

and as a result we obtain the volume formula

$$\text{vol}(\mathcal{D}_{\theta_{\max}}) = (2\theta_{\max})^m. \quad (45)$$

Alternatively, we may also consider the mapping

$$\varphi: \theta = [\theta_0 \quad \cdots \quad \theta_{m-1}] \mapsto \sum_{x=0}^{m-1} e^{i\theta_x} |x\rangle\langle x| \quad (46)$$

that parameterizes all diagonal unitaries with phase angles $-\pi \leq \theta_x < \pi$. This is a bijection, and we may define the pushforward measure on the set of diagonal unitaries using the Lebesgue measure λ as

$$\varphi_*(\lambda)(\mathcal{B}) = \lambda(\varphi^{-1}(\mathcal{B})). \quad (47)$$

The measure $\varphi_*(\lambda)$ can be verified to be invariant under multiplication and, up to a constant factor, is the unique Haar measure for the group of diagonal unitaries. This then gives the volume $\varphi_*(\lambda)(\mathcal{D}_{\theta_{\max}}) = (2\theta_{\max})^m$ same as our previous formula.

3 Faster Trotter steps using block encoding

We now introduce our first method to perform faster Trotter steps based on block encoding. The structure of this section is as follows. We describe how to reduce the 1-norm of Hamiltonian terms using product formulas in [Section 3.1](#), removing the bottleneck of the qubitization algorithm. We then give a circuit implementation of the preparation and selection subroutines in [Section 3.2](#), assuming that the Hamiltonian coefficients are efficiently computable. This gives simulations with complexity nearly linear in the spacetime volume for power-law exponent $\alpha \geq 2$. Readers who wish to see the overall algorithm may skip ahead to [Section 3.3](#).

3.1 Reducing the 1-norm of Hamiltonian coefficients

Consider an n -qubit Hamiltonian $H = \sum_{1 \leq j < k \leq n} H_{j,k}$, where $H_{j,k}$ are 2-local terms acting non-trivially only on sites j and k with norm $\|H_{j,k}\| \leq 1/|j-k|^\alpha$. Our aim is to implement Trotter steps using block encodings and qubitization. To this end, we rewrite the Hamiltonian in the Pauli basis:

$$\begin{aligned}
H = & \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(xx)} X_j X_k + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(xy)} X_j Y_k + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(xz)} X_j Z_k \\
& + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(yx)} Y_j X_k + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(yy)} Y_j Y_k + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(yz)} Y_j Z_k \\
& + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(zx)} Z_j X_k + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(zy)} Z_j Y_k + \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(zz)} Z_j Z_k \\
& + \sum_{1 \leq j \leq n} \beta_j^{(x)} X_j + \sum_{1 \leq j \leq n} \beta_j^{(y)} Y_j + \sum_{1 \leq j \leq n} \beta_j^{(z)} Z_j + \beta^{(i)} I,
\end{aligned} \tag{48}$$

where the coefficients have the scaling

$$\beta_{j,k}^{(\sigma,\sigma')} = \mathcal{O}\left(\frac{1}{|j-k|^\alpha}\right), \quad \beta_j^{(\sigma)} = \begin{cases} \mathcal{O}(1), & \alpha > 1, \\ \mathcal{O}(\log n), & \alpha = 1, \\ \mathcal{O}(n^{1-\alpha}), & 0 < \alpha < 1, \end{cases} \quad \beta^{(i)} = \begin{cases} \mathcal{O}(n), & \alpha > 1, \\ \mathcal{O}(n \log n), & \alpha = 1, \\ \mathcal{O}(n^{2-\alpha}), & 0 < \alpha < 1, \end{cases} \tag{49}$$

for $\sigma \neq \sigma' \in \{x, y, z\}$. Since $\beta^{(i)} I$ commutes with all the other terms and its evolution only introduces a global phase, we may separate out this term without introducing error. We then use product formulas to make a coarse-grained decomposition of the remaining evolution into exponentials of 12 groups of terms. The commutator norm Eq. (28) corresponding to this decomposition has the scaling [32, Theorem H.2]

$$\tilde{\alpha}_{\text{comm}}^{o(1)} = \begin{cases} n^{o(1)}, & \alpha \geq 1, \\ n^{1-\alpha+o(1)}, & 0 < \alpha < 1, \end{cases} \tag{50}$$

which implies

$$r = \begin{cases} \frac{n^{o(1)} t^{1+o(1)}}{e^{o(1)}}, & \alpha \geq 1, \\ \frac{n^{1-\alpha+o(1)} t^{1+o(1)}}{e^{o(1)}}, & 0 < \alpha < 1. \end{cases} \tag{51}$$

In particular, the Trotter error bound corresponding to this coarse-grained decomposition is never larger than that of the naive decomposition (in which all $\Theta(n^2)$ terms are splitted).

Note that all on-site terms from the same group pairwise commute and can be exponentiated with complexity $\mathcal{O}(n)$. So without loss of generality, we will only consider Hamiltonians of the form⁴

$$H = \sum_{1 \leq j < k \leq n} \beta_{j,k} X_j Y_k. \tag{52}$$

The coefficients of this Hamiltonian have vector 1-norm scaling like [32, Lemma H.1]

$$\|\beta\|_1 = \sum_{1 \leq j < k \leq n} |\beta_{j,k}| = \begin{cases} \mathcal{O}(n), & \alpha > 1, \\ \mathcal{O}(n \log n), & \alpha = 1, \\ \mathcal{O}(n^{2-\alpha}), & 0 < \alpha < 1. \end{cases} \tag{53}$$

⁴Note that H is indeed Hermitian as X_j and Y_k commute for $j < k$.

For $\alpha > 1$, even if we use an improved qubitization circuit with product structure like in Eq. (39), we still need a cost roughly

$$\underbrace{n}_{\text{block encoding}} \cdot \underbrace{n \left(\frac{t}{r}\right)^{o(1)}}_{\text{effective time}} = \underbrace{n^2 \left(\frac{t}{r}\right)^{o(1)}}_{\text{Trotter step complexity}} \quad (54)$$

to simulate for an almost constant time, to implement a single Trotter step. Thus, block encoding does not seem to offer a significant benefit over the naive approach. In what follows, we show that product formulas can be used to reduce the 1-norm of Hamiltonian coefficients, which resolves this technical obstacle and, when combined with block encoding, provide a fast implementation of Trotter steps.

To simplify the discussion, we assume that the system size n is a power of 2. We also use

$$\begin{aligned} H_{[j,k]} &:= \sum_{j \leq u < v \leq k} \beta_{u,v} X_u Y_v \quad (1 \leq j < k \leq n), \\ H_{[j,k]:[l,m]} &:= \sum_{\substack{j \leq u \leq k \\ l \leq v \leq m}} \beta_{u,v} X_u Y_v \quad (1 \leq j \leq k < l \leq m \leq n) \end{aligned} \quad (55)$$

to represent terms within a specific interval and across two disjoint intervals of sites. Then, we define a decomposition via the recurrence relation

$$H_{[j,k]} = H_{[j, \lfloor \frac{j+k}{2} \rfloor]: [\lfloor \frac{j+k}{2} \rfloor + 1, k]} + H_{[j, \lfloor \frac{j+k}{2} \rfloor]} + H_{[\lfloor \frac{j+k}{2} \rfloor + 1, k]}, \quad (56)$$

which unwraps to

$$\begin{aligned} H &= H_{[1,n]} \\ &= H_{[1, \frac{n}{2}]: [\frac{n}{2} + 1, n]} + H_{[1, \frac{n}{2}]} + H_{[\frac{n}{2} + 1, n]} \\ &= H_{[1, \frac{n}{2}]: [\frac{n}{2} + 1, n]} + H_{[1, \frac{n}{4}]: [\frac{n}{4} + 1, \frac{n}{2}]} + H_{[\frac{n}{2} + 1, \frac{3n}{4}]: [\frac{3n}{4} + 1, n]} + H_{[1, \frac{n}{4}]} + H_{[\frac{n}{4} + 1, \frac{n}{2}]} + H_{[\frac{n}{2} + 1, \frac{3n}{4}]} + H_{[\frac{3n}{4} + 1, n]} \\ &= \dots \\ &= \sum_{\ell=1}^{\log n - 1} \sum_{b=0}^{2^{\ell-1} - 1} H_{[2b \frac{n}{2^\ell} + 1, (2b+1) \frac{n}{2^\ell}]: [(2b+1) \frac{n}{2^\ell} + 1, 2(b+1) \frac{n}{2^\ell}]}. \end{aligned} \quad (57)$$

Here, the variable ℓ labels the layer of the decomposition and b indexes the pairs of neighboring intervals within layer ℓ . See Figure 1 for an illustration of this recursion for $\ell = 1, 2, 3$. Correspondingly, we apply product formulas to decompose the entire evolution into exponentials of each individual $H_{[2b \frac{n}{2^\ell} + 1, (2b+1) \frac{n}{2^\ell}]: [(2b+1) \frac{n}{2^\ell} + 1, 2(b+1) \frac{n}{2^\ell}]}$. Similar to the above analysis, this coarse-grained decomposition gives a Trotter error bound no larger than that of the fine-grained decomposition and implies the scaling Eq. (51) as well.

To proceed, we take a closer look at this recursive decomposition and collect some of its features below.

- (i) There are $\log n - 1 = \mathcal{O}(\log n)$ layers in the decomposition, all indexed by ℓ .
- (ii) For a fixed layer ℓ , there are 2^ℓ consecutive intervals each of length $\frac{n}{2^\ell}$.
- (iii) Within each layer ℓ , even and odd intervals are further grouped into $2^\ell - 1$ pairs (indexed by b). Terms across intervals from the same pair are denoted by $H_{[2b \frac{n}{2^\ell} + 1, (2b+1) \frac{n}{2^\ell}]: [(2b+1) \frac{n}{2^\ell} + 1, 2(b+1) \frac{n}{2^\ell}]}$. The total number of pairs of intervals is $\mathcal{O}(n)$ by the master theorem Lemma 1.




Figure 1: Illustration of the recursive decomposition of $\sum_{j < k} \beta_{j,k} X_j Y_k$ with 3 layers. The black dots represent qubits and the curves correspond to the interactions between them. In each layer, we cut the system into equal halves (dashed lines) and consider only the interactions across the cuts. The colors indicate whether the interaction consists of Pauli- X (shaded orange) or Pauli- Y (green).

- (iv) When $\ell = 1, \dots, \log n - 1$ and $b = 0, \dots, 2^{\ell-1} - 1$, $H_{[2b \frac{n}{2^\ell} + 1, (2b+1) \frac{n}{2^\ell}]: [(2b+1) \frac{n}{2^\ell} + 1, 2(b+1) \frac{n}{2^\ell}]}$ provide a partition of all the terms in the original Hamiltonian H .

We may use these features to further simplify our discussion. For instance, since we only have 2-local terms of the same Pauli-type acting across disjoint intervals, we can simultaneously change the basis and consider only Pauli- Z interactions without loss of generality:

$$H_{[2b \frac{n}{2^\ell} + 1, (2b+1) \frac{n}{2^\ell}]: [(2b+1) \frac{n}{2^\ell} + 1, 2(b+1) \frac{n}{2^\ell}]} = \sum_{\substack{2b \frac{n}{2^\ell} + 1 \leq u \leq (2b+1) \frac{n}{2^\ell} \\ (2b+1) \frac{n}{2^\ell} + 1 \leq v \leq 2(b+1) \frac{n}{2^\ell}}} \beta_{u,v} Z_u Z_v. \quad (58)$$

Also, due to the nature of the recursive decomposition, it suffices to focus on implementing a specific $H_{[2b \frac{n}{2^\ell} + 1, (2b+1) \frac{n}{2^\ell}]: [(2b+1) \frac{n}{2^\ell} + 1, 2(b+1) \frac{n}{2^\ell}]}$: we will have similar complexities $\text{cost}(\cdot)$ for all pairs of intervals from different layers, so the total complexity

$$\text{cost}_{\text{rec}}(n) = \sum_{\ell=1}^{\log n - 1} 2^{\ell-1} \text{cost}\left(\frac{n}{2^{\ell-1}}\right) \quad (59)$$

can be immediately bounded by the master theorem [Lemma 1](#). For simplicity, we choose $\ell = 1$, $b = 0$ and we study the complexity $\text{cost}(n)$ of exponentiating $H_{[1, \frac{n}{2}]: [\frac{n}{2} + 1, n]}$.

We note that prior work has utilized above features of the recursive decomposition to develop measurement schedules for local quantum observables [\[18\]](#). For the purpose of quantum simulation, we will however need an additional key observation.

- (v) Coefficients of $H_{[1, \frac{n}{2}]: [\frac{n}{2} + 1, n]}$ have the vector 1-norm scaling:

$$\|\beta\|_1 = \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2} + 1 \leq v \leq n}} |\beta_{u,v}| = \begin{cases} \mathcal{O}(1), & \alpha > 2, \\ \mathcal{O}(\log n), & \alpha = 2, \\ \mathcal{O}(n^{2-\alpha}), & 0 < \alpha < 2. \end{cases} \quad (60)$$

Comparing Eq. [\(60\)](#) with Eq. [\(53\)](#), the coefficients of $H_{[1, \frac{n}{2}]: [\frac{n}{2} + 1, n]}$ have a significantly smaller 1-norm than that of the full Hamiltonian H . In particular, the 1-norm in Eq. [\(60\)](#) is asymptotically

constant for $\alpha > 2$, a factor of n smaller than that of the original Hamiltonian. Therefore, exponentiating $H_{[1, \frac{n}{2}]:[\frac{n}{2}+1, n]}$ would also take a factor of n fewer quantum gates than exponentiating H . Meanwhile, the gate complexity of block encoding remains asymptotically the same: we will show in [Section 3.2](#) that block encoding can be implemented for the pair of intervals $[1, \frac{n}{2}]$ and $[\frac{n}{2} + 1, n]$ with $\text{cost}(n) = \mathcal{O}(n \text{polylog}(nt/\epsilon))$, which gives the total cost by applying the master theorem to Eq. (59). This leads to the desired faster Trotter steps for simulating power-law Hamiltonians.

3.2 Preparation and selection subroutines

In the previous subsection, we have reduced the 1-norm of power-law Hamiltonians by applying a recursive decomposition using product formulas. For completeness, we now give an explicit circuit for block encoding and apply qubitization to simulate the decomposed Hamiltonian.

To be precise, our goal is to simulate the following Hamiltonian acting across the intervals $[1, \frac{n}{2}]$ and $[\frac{n}{2} + 1, n]$:

$$H_{[1, \frac{n}{2}]:[\frac{n}{2}+1, n]} = \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} \beta_{u,v} Z_u Z_v, \quad (61)$$

where the coefficients $|\beta_{u,v}| \leq 1/|u - v|^\alpha$. Additionally, we assume that the coefficients are *logarithmically computable*, meaning the oracle

$$O_\beta |u, v, 0\rangle = |u, v, \beta_{u,v}\rangle \quad (62)$$

can be implemented with cost $\mathcal{O}(\text{polylog}(nt/\epsilon))$.⁵ Note that by making this assumption, we have implicitly assumed certain structural properties of the Hamiltonian coefficients. For example, when $\beta_{u,v} = 1/|u - v|^\alpha$ decays exactly as a power law, the oracle O_β can be implemented with cost $\mathcal{O}(\text{polylog}(nt/\epsilon))$, fulfilling this requirement. However, when Hamiltonian coefficients are arbitrarily given, we prove a lower bound in [Appendix A](#) showing that at least $\sim n^2$ gates are required to implement O_β in the circuit model.

In what follows, we analyze the number of queries to O_β as well as additional inter-query gates. We aim to block-encode $H_{[1, \frac{n}{2}]:[\frac{n}{2}+1, n]}$ with its 1-norm scaling close to Eq. (60). As suggested in [Section 2.3](#), we may simply choose the selection subroutine to have the product structure

$$\sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} |uv\rangle\langle uv| \otimes Z_u Z_v. \quad (63)$$

However, some extra efforts are needed to define the preparation subroutine. In particular, the following black-box state preparation subroutine

$$\begin{aligned} \frac{1}{n/2} \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} |u\rangle|v\rangle &\xrightarrow{O_\beta} \frac{1}{n/2} \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} |u\rangle|v\rangle|\beta_{u,v}\rangle \\ &\mapsto \frac{1}{n/2} \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} |u\rangle|v\rangle|\beta_{u,v}\rangle \left(\sqrt{|\beta_{u,v}|} |0\rangle + \sqrt{1 - |\beta_{u,v}|} |1\rangle \right) \\ &\xrightarrow{O_\beta^\dagger} \frac{1}{n/2} \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} \left(\sqrt{|\beta_{u,v}|} |u\rangle|v\rangle|0\rangle + \sqrt{1 - |\beta_{u,v}|} |\psi_\perp\rangle|1\rangle \right) \end{aligned} \quad (64)$$

⁵This may be relaxed to require the coefficients $\beta_{u,v}$ are *linearly computable* (with cost $\mathcal{O}(n \text{polylog}(nt/\epsilon))$ to perform O_β).

does not work since it enlarges the 1-norm of block encoding by a factor of $\frac{n}{2}$, preventing us to achieve the scaling in Eq. (60).

The issue with the naive preparation subroutine is that we have prepared uniform superposition states at the beginning which ignored the power-law decaying pattern of $\beta_{u,v}$. To address this issue, we group the coefficients into multiple boxes. We take the number of boxes to be logarithmic in the system size n so we can prepare superposition states over these boxes efficiently. Meanwhile, we ensure that the coefficients within each box are approximately uniform so we can again prepare the corresponding superposition state efficiently. This idea was previously used for the efficient block encoding of the plane-wave-basis electronic structure Hamiltonian in first quantization [7]. Here, we describe a variant of this technique to efficiently block-encode power-law Hamiltonians.

To simplify the notation, we shift the lattice and consider the rectangle region $(-\frac{n}{2}, -1] \times [1, \frac{n}{2})$. We then divide it into boxes

$$\mathcal{B}_{\mu,\nu} := \{(u, v) \mid -2^{\mu+1} < u \leq -2^\mu, 2^\nu \leq v < 2^{\nu+1}\} \quad (65)$$

for $\mu, \nu = 0, 1, \dots, \log n - 2$. With the restricted max-norm

$$\|\beta\|_{\max, \mathcal{B}_{\mu,\nu}} = \max_{\substack{-2^{\mu+1} < u \leq -2^\mu \\ 2^\nu \leq v < 2^{\nu+1}}} |\beta_{u,v}|, \quad (66)$$

we will describe a state preparation subroutine whose complexity depends on the ratio between $\|\beta\|_1$ and

$$\|\beta\|_{1, \text{box}} := \sum_{\mu, \nu=0}^{\log n - 2} 2^{\mu+\nu} \|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}. \quad (67)$$

Thus, our approach will be efficient if the coefficients are almost uniformly distributed within each box. The boundary terms can be handled separately without changing the asymptotic complexity scaling.

We start by preparing the state

$$\frac{1}{\sqrt{\|\beta\|_{1, \text{box}}}} \sum_{\mu, \nu=0}^{\log n - 2} \sqrt{2^{\mu+\nu} \|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}} |e_\mu\rangle |e_\nu\rangle, \quad (68)$$

where we use unary encoding to represent μ and ν . This state can be prepared using $\mathcal{O}(\log^2 n)$ gates. Conditioned on the unary value of μ and ν , we prepare the uniform superposition

$$\frac{1}{\sqrt{\|\beta\|_{1, \text{box}}}} \sum_{\mu, \nu=0}^{\log n - 2} \sqrt{2^{\mu+\nu} \|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}} |e_\mu\rangle |e_\nu\rangle \frac{1}{\sqrt{2^\mu}} \sum_{-2^{\mu+1} < u \leq -2^\mu} |u\rangle \frac{1}{\sqrt{2^\nu}} \sum_{2^\nu \leq v < 2^{\nu+1}} |v\rangle, \quad (69)$$

where we represent u and v in binary. We can achieve this by first transforming $|e_\mu\rangle$ and $|e_\nu\rangle$ into run-length unary representation, and performing a sequence of $\mathcal{O}(\log n)$ controlled Hadamard gates to generate the superposition. We now introduce an auxiliary uniform superposition state

$$\frac{1}{\sqrt{\|\beta\|_{1, \text{box}}}} \sum_{\mu, \nu=0}^{\log n - 2} \sqrt{2^{\mu+\nu} \|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}} |e_\mu\rangle |e_\nu\rangle \frac{1}{\sqrt{2^{\mu+\nu}}} \sum_{\substack{-2^{\mu+1} < u \leq -2^\mu \\ 2^\nu \leq v < 2^{\nu+1}}} |u\rangle |v\rangle \frac{1}{\sqrt{\Xi}} \sum_{\omega=0}^{\Xi-1} |\xi\rangle \quad (70)$$

and test the inequality

$$\frac{\xi}{\Xi} < \frac{|\beta_{u,v}|}{\|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}}. \quad (71)$$

This can be performed with one query to the coefficient oracle O_β and mild uses of other gates as follows. We first apply O_β to load the coefficients $\beta_{u,v}$. We also load the values of $\|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}$ via the transformation

$$|e_\mu, e_\nu, 0\rangle \mapsto |e_\mu, e_\nu, \|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}\rangle. \quad (72)$$

Note that there are only $\mathcal{O}(\log^2 n)$ different values of $\|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}$, each of which can be represented using at most $\mathcal{O}(\log(nt/\epsilon))$ bits, so we can efficiently perform this transformation. The inequality test can then be performed with a complexity scaling with the maximum size of the inputs. Assuming Ξ is sufficiently large, the state after the inequality test (omitting the failure part and garbage registers) becomes

$$\begin{aligned} & \frac{1}{\sqrt{\|\beta\|_{1,\text{box}}}} \sum_{\mu,\nu=0}^{\log n-2} \sqrt{2^{\mu+\nu} \|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}} |e_\mu\rangle |e_\nu\rangle \frac{1}{\sqrt{2^{\mu+\nu}}} \sum_{\substack{-2^{\mu+1} < u \leq -2^\mu \\ 2^\nu \leq v < 2^{\nu+1}}} \sqrt{\frac{|\beta_{u,v}|}{\|\beta\|_{\max, \mathcal{B}_{\mu,\nu}}}} |u\rangle |v\rangle \\ &= \frac{1}{\sqrt{\|\beta\|_{1,\text{box}}}} \sum_{\mu,\nu=0}^{\log n-2} \sum_{\substack{-2^{\mu+1} < u \leq -2^\mu \\ 2^\nu \leq v < 2^{\nu+1}}} \sqrt{|\beta_{u,v}|} |e_\mu\rangle |e_\nu\rangle |u\rangle |v\rangle. \end{aligned} \quad (73)$$

In practice, we choose a finite value of Ξ which results in an erroneous block encoding that can be analyzed as in our second remark in [Section 2.3](#). To achieve an overall error of at most ϵ , it suffices to set $\Xi = \text{poly}(nt/\epsilon)$. Under this nested-boxes representation, the selection subroutine becomes

$$\sum_{\mu,\nu=0}^{\log n-2} \sum_{\substack{-2^{\mu+1} < u \leq -2^\mu \\ 2^\nu \leq v < 2^{\nu+1}}} |e_\mu e_\nu\rangle \langle e_\mu e_\nu| \otimes |uv\rangle \langle uv| \otimes Z_{\mu,u} Z_{\nu,v}. \quad (74)$$

Again, this operation has a product structure and can be implemented with gate complexity $\mathcal{O}(n \log n)$ per [Section 2.3](#).

It is clear that we have prepared a state proportional to the desired state for block encoding when all $\beta_{u,v} > 0$. If $\beta_{u,v} < 0$, we simply introduce a minus sign in the implementation of the selection subroutine. The success probability is given by

$$\frac{\sum_{\mu,\nu=0}^{\log n-2} \sum_{\substack{-2^{\mu+1} < u \leq -2^\mu \\ 2^\nu \leq v < 2^{\nu+1}}} |\beta_{u,v}|}{\|\beta\|_{1,\text{box}}} = \frac{\|\beta\|_1}{\|\beta\|_{1,\text{box}}}. \quad (75)$$

To boost this success probability to close to 1, we can perform $\mathcal{O}(\sqrt{\|\beta\|_{1,\text{box}} / \|\beta\|_1})$ steps of amplitude amplification. In other words, the complexity of state preparation will indeed depend on the ratio between $\|\beta\|_{1,\text{box}}$ and $\|\beta\|_1$ as previously claimed.

If the distribution of Hamiltonian coefficients exactly matches a power law, we have

$$\|\beta\|_1 = \sum_{\substack{-\frac{n}{2} < u \leq -1 \\ 1 \leq v < \frac{n}{2}}} |\beta_{u,v}| = \begin{cases} \Theta(1), & \alpha > 2, \\ \Theta(\log n), & \alpha = 2, \\ \Theta(n^{2-\alpha}), & 0 < \alpha < 2. \end{cases} \quad (76)$$

Furthermore, $\|\beta\|_1$ differs from $\|\beta\|_{1,\text{box}}$ by at most a constant factor. This is because

$$\begin{aligned} \sum_{\mu,\nu=0}^{\log n-2} 2^{\mu+\nu} \min_{\substack{-2^{\mu+1}<u\leq-2^\mu \\ 2^\nu\leq v<2^{\nu+1}}} |\beta_{u,v}| &< \sum_{\mu,\nu=0}^{\log n-2} \sum_{\substack{-2^{\mu+1}<u\leq-2^\mu \\ 2^\nu\leq v<2^{\nu+1}}} |\beta_{u,v}| = \|\beta\|_1 \\ &< \|\beta\|_{1,\text{box}} = \sum_{\mu,\nu=0}^{\log n-2} 2^{\mu+\nu} \max_{\substack{-2^{\mu+1}<u\leq-2^\mu \\ 2^\nu\leq v<2^{\nu+1}}} |\beta_{u,v}|, \end{aligned} \quad (77)$$

where the first and last quantities differ termwise by at most a factor of $2^\alpha = \mathcal{O}(1)$. Therefore, the number of amplitude amplification steps is constant. The cost of the preparation subroutine is then dominated by the query to the oracle O_β , which costs $\mathcal{O}(\text{polylog}(nt/\epsilon))$ to implement by our assumption.

We now use qubitization (Lemma 3) to simulate $H_{[1,\frac{n}{2}]:[\frac{n}{2}+1,n]}$ for time t/r with accuracy $\mathcal{O}(\epsilon/(rn))$, where

$$r = \begin{cases} \frac{n^{\alpha(1)} t^{1+\alpha(1)}}{e^{\alpha(1)}}, & \alpha \geq 1, \\ \frac{n^{1-\alpha+\alpha(1)} t^{1+\alpha(1)}}{e^{\alpha(1)}}, & 0 < \alpha < 1, \end{cases} \quad (78)$$

with the selection and preparation subroutine defined above. We bound the gate complexity $\text{cost}(n)$ as follows:

$$\begin{aligned} \text{cost}(n) &= \underbrace{\mathcal{O}\left(\|\beta\|_1 \frac{t}{r} + \log\left(\frac{rn}{\epsilon}\right)\right)}_{\text{number of qubitization steps}} \cdot \left(\underbrace{\mathcal{O}(n \log n)}_{\text{selection}} + \underbrace{\mathcal{O}\left(\sqrt{\frac{\|\beta\|_{1,\text{box}}}{\|\beta\|_1}} n \text{polylog}(nt/\epsilon)\right)}_{\text{preparation}} \right) \\ &= \begin{cases} \mathcal{O}\left(\left(\frac{t}{r} + 1\right) n \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & \alpha \geq 2, \\ \mathcal{O}\left(\left(n^{2-\alpha} \frac{t}{r} + 1\right) n \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & 0 < \alpha < 2. \end{cases} \end{aligned} \quad (79)$$

By using the same evolution and target accuracy with different values of n and invoking the master theorem (Lemma 1), we obtain the gate complexity of implementing one Trotter step

$$\text{cost}_{\text{rec}}(n) = \sum_{\ell=1}^{\log n-1} 2^{\ell-1} \text{cost}\left(\frac{n}{2^{\ell-1}}\right) = \begin{cases} \mathcal{O}\left(\left(\frac{t}{r} + 1\right) n \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & \alpha \geq 2, \\ \mathcal{O}\left(\left(n^{2-\alpha} \frac{t}{r} + 1\right) n \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & 0 < \alpha < 2. \end{cases} \quad (80)$$

This then gives the complexity of the entire quantum simulation:

$$\begin{cases} nt \left(\frac{nt}{\epsilon}\right)^{\alpha(1)}, & \alpha \geq 2, \\ n^{3-\alpha} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + nt \left(\frac{nt}{\epsilon}\right)^{\alpha(1)}, & 1 \leq \alpha < 2, \\ n^{3-\alpha} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + n^{2-\alpha} t \left(\frac{nt}{\epsilon}\right)^{\alpha(1)}, & 0 < \alpha < 1. \end{cases} \quad (81)$$

In general, the complexity of block encoding will depend on the closeness of distribution of Hamiltonian coefficients to a power-law distribution. To quantify this, we let λ_{block} be the maximum ratio between $\|\beta\|_{1,\text{box}}$ and $\|\beta\|_1$, maximized over all pairs of intervals in the hierarchical decomposition Eq. (57). More explicitly,

$$\lambda_{\text{block}} := \max_{\sigma \neq \sigma' \in \{x,y,z\}} \max_{\ell=1,\dots,\log n-1} \max_{b=0,\dots,2^{\ell-1}-1} \frac{\left\| \beta(\sigma,\sigma') \right\|_{1,\text{box},\mathcal{I}_{\ell,b}}}{\left\| \beta(\sigma,\sigma') \right\|_{1,\mathcal{I}_{\ell,b}}}, \quad (82)$$

where $\|\cdot\|_{1, \mathcal{I}_{\ell, b}}$ and $\|\cdot\|_{1, \text{box}, \mathcal{I}_{\ell, b}}$ are the above norms defined with respect to the region

$$\mathcal{I}_{\ell, b} := \left\{ (u, v) \mid 2b \frac{n}{2^\ell} + 1 \leq u \leq (2b+1) \frac{n}{2^\ell}, (2b+1) \frac{n}{2^\ell} + 1 \leq v \leq 2(b+1) \frac{n}{2^\ell} \right\}. \quad (83)$$

Then, the complexity of quantum simulation should be revised to

$$\begin{cases} (\sqrt{\lambda_{\text{block}}} + n) t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & \alpha \geq 2, \\ (\sqrt{\lambda_{\text{block}}} + n) \left(n^{2-\alpha} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + t \left(\frac{nt}{\epsilon}\right)^{o(1)} \right), & 1 \leq \alpha < 2, \\ (\sqrt{\lambda_{\text{block}}} + n) \left(n^{2-\alpha} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + n^{1-\alpha} t \left(\frac{nt}{\epsilon}\right)^{o(1)} \right), & 0 < \alpha < 1. \end{cases} \quad (84)$$

It is clear that when the distribution of coefficients is close to a power law, Eq. (77) implies that $\lambda_{\text{block}} = \mathcal{O}(1)$ so we have recovered the cost scaling claimed earlier. When coefficients deviate significantly from a power law distribution, our gate complexity will enlarge by a factor of $\sqrt{\lambda_{\text{block}}}$ due to the use of amplitude amplification. However, one could then adjust the state preparation subroutine to accommodate the new distribution of coefficients. We leave a systematic study of the general case for future work.

3.3 Summary of the algorithm

We now summarize the block-encoding method for implementing faster Trotter steps:

1. Construct a pre-amplified preparation subroutine for Eq. (73).
2. Perform $\mathcal{O}(\sqrt{\lambda_{\text{block}}})$ steps of amplitude amplification to construct the actual preparation subroutine with λ_{block} defined in Eq. (82).
3. Define the selection subroutine according to Eq. (74).
4. Use qubitization (Lemma 3) to simulate for time t/r with accuracy $\mathcal{O}(\epsilon/(rn))$, with r scaling like Eq. (51).
5. Perform qubitization for all combinations of Pauli operators $\sigma, \sigma' = x, y, z$, layers indexed by $\ell = 1, \dots, \log n - 1$ and pairs of intervals indexed by $b = 0, \dots, 2^{\ell-1} - 1$, to implement a single Trotter step.
6. Repeat r Trotter steps to simulate the entire evolution.

Theorem 1 (Faster Trotter steps using block encoding). *Consider 2-local Hamiltonians $H = \sum_{\sigma \neq \sigma' \in \{x, y, z\}} \sum_{1 \leq j < k \leq n} \beta_{j, k}^{(\sigma, \sigma')} P_j^{(\sigma)} P_k^{(\sigma')}$, where $|\beta_{j, k}^{(\sigma, \sigma')}| \leq 1/|j - k|^\alpha$ for some constant $\alpha > 0$ and $P^{(\sigma)}$ ($\sigma = x, y, z$) are Pauli matrices. Let $t > 0$ be the simulation time and $\epsilon > 0$ be the target accuracy. Assume that the coefficient oracle*

$$O_\beta |\sigma, \sigma', j, k, 0\rangle = |\sigma, \sigma', j, k, \beta_{j, k}^{(\sigma, \sigma')}\rangle \quad (85)$$

can be implemented with gate complexity $\mathcal{O}(\text{polylog}(nt/\epsilon))$. Then H can be simulated with complexity

$$\begin{cases} (\sqrt{\lambda_{\text{block}}} + n) t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & \alpha \geq 2, \\ (\sqrt{\lambda_{\text{block}}} + n) \left(n^{2-\alpha} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + t \left(\frac{nt}{\epsilon}\right)^{o(1)} \right), & 1 \leq \alpha < 2, \\ (\sqrt{\lambda_{\text{block}}} + n) \left(n^{2-\alpha} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + n^{1-\alpha} t \left(\frac{nt}{\epsilon}\right)^{o(1)} \right), & 0 < \alpha < 1. \end{cases} \quad (86)$$

Here, $\lambda_{\text{block}} \geq 1$ defined in Eq. (82) quantifies the deviation of the Hamiltonian coefficient distribution from a power law ($\lambda_{\text{block}} = \mathcal{O}(1)$ if the coefficient distribution exactly matches a power law).

4 Faster Trotter steps using average-cost simulation

In the previous section, we have described a block-encoding-based method to simulate power-law Hamiltonians with efficiently computable coefficients. The complexity of our method is almost linear in the spacetime volume $(nt)^{1+o(1)}$ for $\alpha \geq 2$ and is close to $n^{3-\alpha+o(1)}t^{1+o(1)}$ for $\alpha < 2$, both of which improve the best results from previous work.

In this section, we obtain a further improvement by performing an average-cost quantum simulation of commuting terms. We explain the basic idea of this technique in Section 4.1, with further details on the circuit implementation presented in Section 4.2. Readers may skip ahead to Section 4.3 for a summary of the entire algorithm.

4.1 Simulating commuting terms with average combination cost

We now show that gate complexities can be further reduced, using the simple fact that commuting Hamiltonian terms can be simulated with an *average-cost linear combination*. To elaborate, consider a Hamiltonian $H = \sum_{\gamma=1}^{\Gamma} H_{\gamma}$ with Hermitian H_{γ} and suppose we have block encodings $G_{\gamma,1}^{\dagger} U_{\gamma} G_{\gamma,0} = H_{\gamma}/\beta_{\gamma}$, where $\beta_{\gamma} > 0$, and $G_{\gamma,0}, G_{\gamma,1} : \mathcal{G} \rightarrow \mathcal{H}$ and $U_{\gamma} : \mathcal{H} \rightarrow \mathcal{H}$ can be implemented with cost $c_{\gamma} > 0$. Then, H can be block-encoded by defining

$$\begin{aligned}
 G_0 : \mathcal{G} &\rightarrow \mathbb{C}^{\Gamma} \otimes \mathcal{H}, & G_0 &= \frac{1}{\sqrt{\|\beta\|_1}} \sum_{\gamma=1}^{\Gamma} \sqrt{\beta_{\gamma}} |\gamma\rangle \otimes G_{\gamma,0} \\
 G_1 : \mathcal{G} &\rightarrow \mathbb{C}^{\Gamma} \otimes \mathcal{H}, & G_1 &= \frac{1}{\sqrt{\|\beta\|_1}} \sum_{\gamma=1}^{\Gamma} \sqrt{\beta_{\gamma}} |\gamma\rangle \otimes G_{\gamma,1} \\
 U : \mathbb{C}^{\Gamma} \otimes \mathcal{H} &\rightarrow \mathbb{C}^{\Gamma} \otimes \mathcal{H}, & U &= \sum_{\gamma=1}^{\Gamma} |\gamma\rangle\langle\gamma| \otimes U_{\gamma}, \\
 G_1^{\dagger} U G_0 : \mathcal{G} &\rightarrow \mathcal{G}, & G_1^{\dagger} U G_0 &= \frac{H}{\|\beta\|_1}.
 \end{aligned} \tag{87}$$

This block encoding has a normalization factor of $\|\beta\|_1$ and can be implemented with cost $\sum_{\gamma} c_{\gamma}$ (plus some additional cost for preparing the ancilla state $\frac{1}{\sqrt{\|\beta\|_1}} \sum_{\gamma=1}^{\Gamma} \sqrt{\beta_{\gamma}} |\gamma\rangle$). Invoking Lemma 3, we can simulate H for time t with accuracy ϵ with a cost scaling like

$$\mathcal{O} \left(\left(\|\beta\|_1 t + \log \left(\frac{1}{\epsilon} \right) \right) \cdot \left(\sum_{\gamma=1}^{\Gamma} c_{\gamma} \right) \right), \tag{88}$$

which reduces to $\sim t \left(\sum_{\gamma} \beta_{\gamma} \right) \left(\sum_{\gamma} c_{\gamma} \right)$ ignoring the error scaling. This is a worst-case combination because we are paying the same total cost $\sum_{\gamma} c_{\gamma}$ for each of the $\sim t \sum_{\gamma} \beta_{\gamma}$ qubitization steps.

This worst-case cost scaling may sometimes be avoided by recursively performing simulation in the interaction picture [81, 87]. Roughly speaking, the recursive interaction-picture approach has a cost scaling like $\sim t \left(\sum_{\gamma} \beta_{\gamma} c_{\gamma} \right) \log^{2\Gamma-1} (t \|\beta\|_{\max})$. Thus we achieve the desired average-case

combination cost, but also pick up a factor that scales exponentially with the number of terms Γ , which prevents the approach from being useful in many cases.

Instead, we make the following simple yet important observation about simulating commuting terms using block encodings and qubitization.

Lemma 4 (Simulating commuting terms with average combination cost). *Let $G_{\gamma,0}, G_{\gamma,1} : \mathcal{G} \rightarrow \mathcal{H}$ be isometries and $U_\gamma : \mathcal{H} \rightarrow \mathcal{H}$ be unitaries such that $G_{\gamma,1}^\dagger U_\gamma G_{\gamma,0} = H_\gamma / \beta_\gamma$ are Hermitian with $\beta_\gamma > 0$. Assume that H_γ pairwise commute. Given a target evolution time t and accuracy ϵ , there exist unitaries $V_{\varphi_\gamma} : \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathcal{H} \rightarrow \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathcal{H}$ parameterized by angles $\varphi_{\gamma,1}, \dots, \varphi_{\gamma,r_\gamma}$ such that*

$$\left\| \prod_{\gamma=1}^{\Gamma} \left(\left(\langle + | \otimes \frac{\langle 0 | \otimes G_{\gamma,0}^\dagger + \langle 1 | \otimes G_{\gamma,1}^\dagger}{\sqrt{2}} \right) V_{\varphi_\gamma} \left(| + \rangle \otimes \frac{| 0 \rangle \otimes G_{\gamma,0} + | 1 \rangle \otimes G_{\gamma,1}}{\sqrt{2}} \right) \right) - e^{-itH} \right\| \leq \epsilon. \quad (89)$$

The number of steps r_γ are even integers with the asymptotic scaling

$$r_\gamma = \mathcal{O} \left(\beta_\gamma t + \log \left(\frac{\Gamma}{\epsilon} \right) \right), \quad (90)$$

and V_{φ_γ} are obtained by applying [Lemma 3](#) to simulate $G_{\gamma,1}^\dagger U_\gamma G_{\gamma,0}$ for time $\beta_\gamma t$ with accuracy ϵ/Γ .

In essence, we are just using the first-order Lie-Trotter formula $e^{-itH} = e^{-itH_\Gamma} \dots e^{-itH_1}$ with each exponential further simulated by the qubitization algorithm. Because Hamiltonian terms pairwise commute, there is no Trotter error introduced in this decomposition. As a result, we can simulate the target Hamiltonian with an average-case combination cost $\sim t \left(\sum_\gamma \beta_\gamma c_\gamma \right)$, without the unwanted exponential factor from the interaction-picture approach.⁶

For power-law Hamiltonians with $\alpha < 2$, recall from [Section 3.1](#) that we can without loss of generality consider

$$H_{(-\frac{n}{2}, -1]:[1, \frac{n}{2}]} = \sum_{\substack{-\frac{n}{2} < u \leq -1 \\ 1 \leq v < \frac{n}{2}}} \beta_{u,v} Z_u Z_v, \quad (91)$$

where we have dropped the subdominant terms and shifted the intervals for notational convenience. Such a Hamiltonian term is directly block-encoded and simulated by qubitization in [Section 3.2](#). We now show how that result can be further improved using the average-cost simulation technique. To this end, we divide each interval into m subintervals at $\pm l_1, \dots, \pm l_m$ and define

$$\mathcal{C}_{j,k} := \{(u, v) \mid -l_{j+1} < u \leq -l_j, l_k \leq v < l_{k+1}\}. \quad (92)$$

See [Figure 2](#) for an illustration of the corresponding decomposition of Hamiltonian. For terms corresponding to $\mathcal{C}_{j,k}$, we have that the 1-norm is asymptotically bounded by

$$\|\beta\|_{1, \mathcal{C}_{j,k}} = \mathcal{O} \left(\frac{(l_{j+1} - l_j)(l_{k+1} - l_k)}{(l_j + l_k)^\alpha} \right). \quad (93)$$

⁶Note that we are using the same qubitization algorithm to simulate all the Hamiltonian terms, so there is a term-independent constant prefactor for all the gate complexities. This is why we have the summation rule $\sum_{\gamma=1}^{\Gamma} c_\gamma \mathcal{O}(\beta_\gamma t + \log(\Gamma/\epsilon)) = \mathcal{O} \left(\left(\sum_{\gamma=1}^{\Gamma} \beta_\gamma c_\gamma \right) t + \sum_{\gamma=1}^{\Gamma} c_\gamma \log(\Gamma/\epsilon) \right)$. This rule will be used without declaration in the remainder of the paper.




Figure 2: Illustration of the interaction pairs in $\mathcal{C}_{j,k}$ defined in Eq. (92). For convenience, we relabel the sites to $-n/2, \dots, -1$ and $1, \dots, n/2$.

We will describe in Section 4.2 how to block-encode the Hamiltonian terms within $\mathcal{C}_{j,k}$ with 1-norm scaling exactly as above. Specifically, we show that up to polylogarithmic factors the preparation and selection subroutines have a total cost of

$$\tilde{\mathcal{O}}(1 + (l_{j+1} - l_j) + (l_{k+1} - l_k)). \quad (94)$$

In our above analysis, we have omitted a factor of λ_{avg} due to the use of amplitude amplification. Just like Section 3.2, this factor is close to 1 when the distribution of Hamiltonian coefficients are close to a power-law distribution. We assume this is the case to simplify the following discussion, and present the full complexity expression in Section 4.2.

We use the uniform division $l_j = \Theta(j \frac{n}{m})$ here for simplicity, although other divisions may lead to circuits with lower cost. Choosing the evolution time t/r and ignoring the $\text{polylog}(nt/\epsilon)$ scaling, we estimate the cost of simulating $H_{(-\frac{n}{2}, -1]:[1, \frac{n}{2}]}$ (for $\alpha < 2$) as

$$\begin{aligned} \tilde{\mathcal{O}} \left(\sum_{j,k=1}^m \left(\frac{\left(\frac{n}{m}\right)^2 \frac{t}{r}}{\left(j \frac{n}{m} + k \frac{n}{m}\right)^\alpha} + 1 \right) \left(1 + \frac{n}{m}\right) \right) &= \tilde{\mathcal{O}} \left(\left(\sum_{j,k=1}^m \frac{\left(\frac{n}{m}\right)^{2-\alpha} \frac{t}{r}}{(j+k)^\alpha} + m^2 \right) \left(1 + \frac{n}{m}\right) \right) \\ &= \tilde{\mathcal{O}} \left(\left(\left(\sum_{j=1}^m \frac{1}{j^{\alpha/2}} \right)^2 \left(\frac{n}{m}\right)^{2-\alpha} \frac{t}{r} + m^2 \right) \left(1 + \frac{n}{m}\right) \right) \\ &= \tilde{\mathcal{O}} \left(\left(\left(m^{1-\frac{\alpha}{2}}\right)^2 \left(\frac{n}{m}\right)^{2-\alpha} \frac{t}{r} + m^2 \right) \left(1 + \frac{n}{m}\right) \right) \\ &= \tilde{\mathcal{O}} \left(\left(n^{2-\alpha} \frac{t}{r} + m^2 \right) \left(1 + \frac{n}{m}\right) \right). \end{aligned} \quad (95)$$

We balance the two scalings in the first parentheses to optimize the gate complexity, which implies

$$m = \begin{cases} \Theta \left(n^{1-\frac{\alpha}{2}} \right), & 1 \leq \alpha < 2, \\ \Theta \left(n^{\frac{1}{2}} \right), & 0 < \alpha < 1. \end{cases} \quad (96)$$

For all values of $\alpha < 2$, one can verify that $m = \Omega(1)$ and $m = \mathcal{O}(n)$ (in fact $m = \mathcal{O}(\sqrt{n})$), so this choice of m is indeed valid. We give a slightly better (yet more complicated) choice of m in Section 4.2.

The remaining analysis proceeds similarly as in Section 3.2. We use qubitization (Lemma 3) to simulate all m^2 pairs of subintervals in $H_{[1, \frac{n}{2}]:[\frac{n}{2}+1, n]}$ for time t/r with accuracy $\mathcal{O}(\epsilon/(rm^2n))$, where

$$r = \begin{cases} \frac{n^{o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}, & \alpha \geq 1, \\ \frac{n^{1-\alpha+o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}, & 0 < \alpha < 1, \end{cases} \quad (97)$$

with an average-case combination cost as discussed above. We have the cost function

$$\begin{cases} \mathcal{O}\left(n^{2-\frac{\alpha}{2}} \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & 1 \leq \alpha < 2, \\ \mathcal{O}\left(n^{\frac{3}{2}} \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & 0 < \alpha < 1. \end{cases} \quad (98)$$

By using the same evolution time and target accuracy with different values of n and invoking the master theorem ([Lemma 1](#)), we obtain the gate complexity of implementing one Trotter step

$$\text{cost}_{\text{rec}}(n) = \sum_{l=1}^{\log n-1} 2^{l-1} \text{cost}\left(\frac{n}{2^{l-1}}\right) = \begin{cases} \mathcal{O}\left(n^{2-\frac{\alpha}{2}} \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & 1 \leq \alpha < 2, \\ \mathcal{O}\left(n^{\frac{3}{2}} \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & 0 < \alpha < 1. \end{cases} \quad (99)$$

This then gives the complexity of the entire quantum simulation:

$$\begin{cases} \frac{n^{2-\alpha/2+o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}, & 1 \leq \alpha < 2, \\ \frac{n^{5/2-\alpha+o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}, & 0 < \alpha < 1. \end{cases} \quad (100)$$

4.2 Preparation and selection subroutines

We now describe a circuit that achieves the average-case combination cost for block encoding power-law Hamiltonians claimed in the previous subsection.

Specifically, let $l_j, l_{j+1}, l_k, l_{k+1}$ be arbitrary integers such that $1 \leq l_j < l_{j+1} \leq \frac{n}{2}$ and $1 \leq l_k < l_{k+1} \leq \frac{n}{2}$. Our goal is to block-encode the Hamiltonian

$$H_{(-l_{j+1}, -l_j]:[l_k, l_{k+1})} = \sum_{\substack{-l_{j+1} < u \leq -l_j \\ l_k \leq v < l_{k+1}}} \beta_{u,v} Z_u Z_v, \quad (101)$$

where $|\beta_{u,v}| \leq 1/|u-v|^\alpha$ and we have shifted the intervals for notational convenience. Additionally, we assume that the coefficients are *logarithmically computable*,⁷ meaning the oracle

$$O_\beta |u, v, 0\rangle = |u, v, \beta_{u,v}\rangle \quad (102)$$

can be implemented with cost $\mathcal{O}(\text{polylog}(nt/\epsilon))$. By making this oracle assumption, we have implicitly assumed certain underlying structure of the Hamiltonian coefficients: in [Appendix A](#) we show that one needs $\sim n^2$ gates to implement O_β in the circuit model when structural properties of coefficients are unavailable. The selection subroutine can simply be chosen as

$$\sum_{-l_{j+1} < u \leq -l_j} \sum_{l_k \leq v < l_{k+1}} |u, v\rangle \langle u, v| \otimes Z_u Z_v. \quad (103)$$

Due to the product structure of this operation, we can implement it with gate complexity $\mathcal{O}\left(\frac{n}{m}\right)$.

In what follows, we analyze the preparation subroutine. Here, we only consider the uniform division $l_j = \Theta\left(j \frac{n}{m}\right)$, as this is enough to justify the gate complexity claimed in [Section 4.1](#). We will use the naive black-box state preparation technique, because coefficients within the divided subintervals are already close to uniform. We start by preparing the uniform superposition

$$\frac{1}{\sqrt{l_{j+1} - l_j}} \sum_{-l_{j+1} < u \leq -l_j} |u\rangle \frac{1}{\sqrt{l_{k+1} - l_k}} \sum_{l_k \leq v < l_{k+1}} |v\rangle. \quad (104)$$

⁷As before, this can be relaxed to linearly computable.

Then we invoke the black-box state preparation subroutine

$$\begin{aligned}
& \frac{1}{\sqrt{(l_{j+1} - l_j)(l_{k+1} - l_k)}} \sum_{\substack{-l_{j+1} < u \leq -l_j \\ l_k \leq v < l_{k+1}}} |u, v\rangle \\
\stackrel{O_\beta}{\mapsto} & \frac{1}{\sqrt{(l_{j+1} - l_j)(l_{k+1} - l_k)}} \sum_{\substack{-l_{j+1} < u \leq -l_j \\ l_k \leq v < l_{k+1}}} |u, v\rangle |\beta_{u,v}\rangle \\
\mapsto & \frac{1}{\sqrt{(l_{j+1} - l_j)(l_{k+1} - l_k)}} \sum_{\substack{-l_{j+1} < u \leq -l_j \\ l_k \leq v < l_{k+1}}} |u, v\rangle |\beta_{u,v}\rangle \left(\sqrt{\frac{|\beta_{u,v}|}{\|\beta\|_{\max}}} |0\rangle + \sqrt{\frac{\|\beta\|_{\max} - |\beta_{u,v}|}{\|\beta\|_{\max}}} |1\rangle \right),
\end{aligned} \tag{105}$$

where the last step can be realized using inequality test in a way similar to [Section 3.2](#). There is no need to uncompute the ancilla register $|\beta_{u,v}\rangle$, as the uncomputation is automatically performed in the qubitization algorithm. It is clear that we have prepared a state proportional to the desired state for block encoding when all $\beta_{u,v} > 0$. If some $\beta_{u,v} < 0$, we simply introduce a minus sign in the implementation of the selection subroutine. The success probability is given by

$$\frac{\sum_{\substack{-l_{j+1} < u \leq -l_j \\ l_k \leq v < l_{k+1}}} |\beta_{u,v}|}{(l_{j+1} - l_j)(l_{k+1} - l_k) \|\beta\|_{\max}} = \frac{\|\beta\|_1}{(l_{j+1} - l_j)(l_{k+1} - l_k) \|\beta\|_{\max}}. \tag{106}$$

To boost this probability to close to 1, we can perform $\mathcal{O}\left(\sqrt{(l_{j+1} - l_j)(l_{k+1} - l_k) \|\beta\|_{\max} / \|\beta\|_1}\right)$ steps of amplitude amplification.

If the distribution of Hamiltonian coefficients exactly matches a power law, we have

$$\|\beta\|_1 = \sum_{\substack{-l_{j+1} < u \leq -l_j \\ l_k \leq v < l_{k+1}}} |\beta_{u,v}| \geq \frac{(l_{j+1} - l_j)(l_{k+1} - l_k)}{(l_{j+1} + l_{k+1})^\alpha}, \quad \|\beta\|_{\max} \leq \frac{1}{(l_j + l_k)^\alpha}, \tag{107}$$

so the number of amplification steps scales like

$$\mathcal{O}\left(\sqrt{\frac{(l_{j+1} - l_j)(l_{k+1} - l_k) \|\beta\|_{\max}}{\|\beta\|_1}}\right) = \mathcal{O}\left(\sqrt{\frac{(l_{j+1} + l_{k+1})^\alpha}{(l_j + l_k)^\alpha}}\right) = \mathcal{O}(1), \tag{108}$$

which justifies the previous claim in Eq. (94). In general, the complexity of block encoding will depend on the closeness of the distribution of Hamiltonian coefficients to a power-law distribution. To quantify this, we let λ_{avg} be the maximum ratio between $(l_{j+1} - l_j)(l_{k+1} - l_k) \|\beta\|_{\max}$ and $\|\beta\|_1$, maximized over all pairs of intervals in the decomposition Eq. (57) with a further uniform division. More explicitly,

$$\begin{aligned}
\lambda_{\text{avg}} := & \max_{\sigma \neq \sigma' \in \{x, y, z\}} \max_{\ell=1, \dots, \log n - 1} \max_{b=0, \dots, 2^\ell - 1} \\
& \frac{(l_{\ell, b, j+1} - l_{\ell, b, j})(l_{\ell, b, k+1} - l_{\ell, b, k}) \left\| \beta^{(\sigma, \sigma')} \right\|_{\max, \mathcal{I}_{\ell, b, j, k}}}{\max_{j, k=1, \dots, m} \left\| \beta^{(\sigma, \sigma')} \right\|_{1, \mathcal{I}_{\ell, b, j, k}}},
\end{aligned} \tag{109}$$

where $l_{\ell, b, j}$ are the uniform division points, and $\|\cdot\|_{\max, \mathcal{I}_{\ell, b, j, k}}$ and $\|\cdot\|_{1, \mathcal{I}_{\ell, b, j, k}}$ are the max- and 1-norm restricted to the region

$$\mathcal{I}_{\ell, b, j, k} := \left\{ (u, v) \mid (2b+1)\frac{n}{2^\ell} - l_{\ell, b, j+1} < u \leq (2b+1)\frac{n}{2^\ell} - l_{\ell, b, j}, \right. \\
\left. (2b+1)\frac{n}{2^\ell} + l_{\ell, b, k} \leq v \leq (2b+1)\frac{n}{2^\ell} + l_{\ell, b, k+1} \right\}. \tag{110}$$

With this definition, the cost of simulating $H_{(-\frac{n}{2}, -1]:[1, \frac{n}{2}]}$ for time t/r should be revised to

$$\text{cost}(n) = \tilde{\mathcal{O}} \left(\left(n^{2-\alpha} \frac{t}{r} + m^2 \right) \left(\sqrt{\lambda_{\text{avg}}} + \frac{n}{m} \right) \right). \quad (111)$$

We balance the first term by choosing

$$m = \max \left\{ \Theta \left(n^{1-\alpha/2} \left(\frac{t}{r} \right)^{1/2} \right), 1 \right\}, \quad (112)$$

which implies through the master theorem ([Lemma 1](#)) that the total simulation has complexity

$$\tilde{\mathcal{O}} \left((n^{2-\alpha} t + r) \left(\sqrt{\lambda_{\text{avg}}} + \min \left\{ n^{\alpha/2} \left(\frac{r}{t} \right)^{1/2}, n \right\} \right) \right). \quad (113)$$

Inserting the scaling of r from Eq. ([51](#)), we obtain

$$\left\{ \begin{array}{l} \min \left\{ \sqrt{\lambda_{\text{avg}}} n^{2-\alpha} t \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} t + n^{2-\frac{\alpha}{2}} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)}, \right. \\ \quad \left. \left(\sqrt{\lambda_{\text{avg}}} n^{2-\alpha} t + n^{3-\alpha} t \right) \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} t + nt \right) \left(\frac{nt}{\epsilon} \right)^{o(1)} \right\}, \quad 1 \leq \alpha < 2, \\ \min \left\{ \sqrt{\lambda_{\text{avg}}} n^{2-\alpha} t \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{1-\alpha} t + n^{\frac{5}{2}-\alpha} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)}, \right. \\ \quad \left. \left(\sqrt{\lambda_{\text{avg}}} n^{2-\alpha} t + n^{3-\alpha} t \right) \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{1-\alpha} t + n^{2-\alpha} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)} \right\}, \quad 0 < \alpha < 1. \end{array} \right. \quad (114)$$

4.3 Summary of the algorithm

We now summarize the average-cost simulation method for implementing faster Trotter steps:

1. Construct a pre-amplified preparation subroutine for Eq. ([105](#)).
2. Perform $\mathcal{O}(\sqrt{\lambda_{\text{avg}}})$ steps of amplitude amplification to construct the actual preparation subroutine with λ_{avg} defined in Eq. ([109](#)).
3. Define the selection subroutine according to Eq. ([103](#)).
4. Use qubitization ([Lemma 3](#)) to simulate for time t/r with accuracy $\mathcal{O}(\epsilon/(rn^2))$, with r scaling like Eq. ([51](#)).
5. Perform qubitization for all combinations of Pauli operators $\sigma, \sigma' = x, y, z$, layers indexed by $\ell = 1, \dots, \log n - 1$, intervals indexed by $b = 0, \dots, 2^{\ell-1} - 1$ and pairs of subintervals indexed by $j, k = 1, \dots, m$, to implement a single Trotter step.
6. Repeat r Trotter steps to simulate the entire evolution.

Theorem 2 (Faster Trotter steps using average-cost simulation). *Consider 2-local Hamiltonians $H = \sum_{\sigma \neq \sigma' \in \{x, y, z\}} \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(\sigma, \sigma')} P_j^{(\sigma)} P_k^{(\sigma')}$, where $|\beta_{j,k}^{(\sigma, \sigma')}| \leq 1/|j-k|^\alpha$ for some constant $\alpha > 0$*

and $P^{(\sigma)}$ ($\sigma = x, y, z$) are Pauli matrices. Let $t > 0$ be the simulation time and $\epsilon > 0$ be the target accuracy. Assume that the coefficient oracle

$$O_\beta |\sigma, \sigma', j, k, 0\rangle = |\sigma, \sigma', j, k, \beta_{j,k}^{(\sigma, \sigma')}\rangle \quad (115)$$

can be implemented with gate complexity $\mathcal{O}(\text{polylog}(nt/\epsilon))$. Then H can be simulated with complexity

$$\left\{ \begin{array}{l} \min \left\{ \sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{2}} t \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{\frac{\alpha}{2}} t + n^{2-\frac{\alpha}{2}} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)}, \right. \\ \left. \left(\sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{2}} t + n^{3-\alpha} t \right) \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{\frac{\alpha}{2}} t + nt \right) \left(\frac{nt}{\epsilon} \right)^{o(1)} \right\}, \quad 1 \leq \alpha < 2, \\ \min \left\{ \sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{2}} t \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{1-\frac{\alpha}{2}} t + n^{\frac{5}{2}-\alpha} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)}, \right. \\ \left. \left(\sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{2}} t + n^{3-\alpha} t \right) \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{1-\frac{\alpha}{2}} t + n^{2-\alpha} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)} \right\}, \quad 0 < \alpha < 1. \end{array} \right. \quad (116)$$

Here, $\lambda_{\text{avg}} \geq 1$ defined in Eq. (109) quantifies the deviation of the Hamiltonian coefficient distribution from a power law ($\lambda_{\text{avg}} = \mathcal{O}(1)$ if the coefficient distribution exactly matches a power law).

5 Faster Trotter steps using low-rank decomposition

In the previous sections, we have showed how the 1-norm of Hamiltonian coefficients can be reduced via a recursive decomposition using product formulas, which results in faster circuit implementation of Trotter steps. We further improve our result by simulating commuting terms with an average-case combination cost. Assuming Hamiltonian coefficients are efficiently computable, these techniques together enable simulations of power-law systems with complexity nearly linear in the spacetime volume for $\alpha \geq 2$, whereas the cost becomes $n^{2-\alpha/2+o(1)} t^{1+o(1)}$ for $1 \leq \alpha < 2$ and $n^{5/2-\alpha+o(1)} t^{1+o(1)}$ for $0 < \alpha < 1$.

In this section, we describe a method for implementing a Trotter step through a recursive decomposition of the Hamiltonian using its hierarchical low-rank structure [46, 50]. This low-rank structure was previously used in [96] to block-encode kernel matrices. Here, we directly use the recursive decomposition in Section 5.1 to construct circuits without block encoding. The overall algorithm and its complexity are then summarized in Section 5.2.

5.1 Recursive low-rank decomposition

The initial steps of our method are the same as that of Section 3.1. In particular, we expand the power-law Hamiltonian in the Pauli basis as in Eq. (48), and use product formulas to perform a coarse-grained decomposition. Without loss of generality, we may focus on 2-local terms

$$H = \sum_{1 \leq j < k \leq n} \beta_{j,k} X_j Y_k, \quad (117)$$

as the on-site terms can be implemented with subdominant cost and the remaining 2-local terms can be handled similarly by a change of basis. As before, we assume that the system size n is a

power of 2 and use

$$\begin{aligned}
H_{[j,k]} &:= \sum_{j \leq u < v \leq k} \beta_{u,v} X_u Y_v \quad (1 \leq j < k \leq n), \\
H_{[j,k]:[l,m]} &:= \sum_{\substack{j \leq u \leq k \\ l \leq v \leq m}} \beta_{u,v} X_u Y_v \quad (1 \leq j \leq k < l \leq m \leq n)
\end{aligned} \tag{118}$$

to represent terms within a specific interval and across two disjoint intervals of sites.

However, we now use a decomposition different from that of [Section 3.1](#). Specifically, we use the recurrence relation

$$\begin{aligned}
H_{[j,k]} &= H_{[j,j+\delta-1]:[j+2\delta,j+3\delta-1]} + H_{[j,j+\delta-1]:[j+3\delta,k]} + H_{[j+\delta,j+2\delta-1]:[j+3\delta,k]} \\
&\quad + H_{[j,j+2\delta-1]} + H_{[j+\delta,j+3\delta-1]} + H_{[j+2\delta,k]} \\
&\quad - H_{[j+\delta,j+2\delta-1]} - H_{[j+2\delta,j+3\delta-1]}
\end{aligned} \tag{119}$$

with $\delta = \lfloor (k - j + 1)/4 \rfloor$, which unwraps to layer $h = \Theta(\log n)$ as

$$\begin{aligned}
H &= H_{[1,n]} \\
&= H_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]} + H_{[1, \frac{n}{4}]:[\frac{3n}{4}+1, n]} + H_{[\frac{n}{4}+1, \frac{n}{2}]:[\frac{3n}{4}+1, n]} \\
&\quad + H_{[1, \frac{n}{2}]} + H_{[\frac{n}{4}+1, \frac{3n}{4}]} + H_{[\frac{n}{2}+1, n]} \\
&\quad - H_{[\frac{n}{4}+1, \frac{n}{2}]} - H_{[\frac{n}{2}+1, \frac{3n}{4}]} \\
&= \dots \\
&= \sum_{\ell=2}^h \sum_{b=0}^{2^{\ell-1}-2} \left(H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+2)\frac{n}{2^\ell}, (2b+3)\frac{n}{2^\ell}]} \right. \\
&\quad + H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]} \\
&\quad \left. + H_{[1+(2b+1)\frac{n}{2^\ell}, (2b+2)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]} \right) \\
&\quad + \sum_{b=0}^{2^h-1} H_{[1+b\frac{n}{2^\ell}, (b+1)\frac{n}{2^\ell}]} + \sum_{b=0}^{2^h-2} H_{[1+b\frac{n}{2^\ell}, (b+1)\frac{n}{2^\ell}]:[1+(b+1)\frac{n}{2^\ell}, (b+2)\frac{n}{2^\ell}]}
\end{aligned} \tag{120}$$

See [Figure 3](#) for an illustration of this decomposition at the first nontrivial layer $\ell = 2$.

We observe the following features of the decomposition that are helpful to describe our circuit implementation:

- (i) There are $h - 1 = \Theta(\log n)$ layers in the decomposition, all indexed by ℓ .
- (ii) For a fixed layer ℓ , there are 2^ℓ consecutive intervals each of length $\frac{n}{2^\ell}$.
- (iii) Within each layer ℓ , intervals are further divided into $2^{\ell-1}$ blocks (indexed by b). Within each pair of consecutive blocks, we only keep the 3 terms that act on intervals with distance at least $\frac{n}{2^\ell}$. They are $H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+2)\frac{n}{2^\ell}, (2b+3)\frac{n}{2^\ell}]}$, $H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]}$, and $H_{[1+(2b+1)\frac{n}{2^\ell}, (2b+2)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]}$. The total number of pairs of intervals is $\mathcal{O}(n)$ by the master theorem [Lemma 1](#).
- (iv) When $\ell = 2, \dots, h$ and $b = 0, \dots, 2^{\ell-1} - 2$, the decomposition provides a partition of all the terms in the Hamiltonian with distance at least $\frac{n}{2^h}$.




Figure 3: Illustration of one iteration [Eq. (119)] in the low-rank decomposition. We repeatedly apply Eq. (119) to the remaining terms in the second row to arrive at the final decomposition Eq. (120).

As in Section 3.1, we may use these features to simplify our analysis. For instance, since we only have 2-local terms with the same type of Pauli operators acting across disjoint intervals, we can simultaneously change the basis and consider only Pauli- Z interactions, e.g.,

$$H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+2)\frac{n}{2^\ell}, (2b+3)\frac{n}{2^\ell}]} = \sum_{\substack{1+2b\frac{n}{2^\ell} \leq u \leq (2b+1)\frac{n}{2^\ell} \\ 1+(2b+2)\frac{n}{2^\ell} \leq v \leq (2b+3)\frac{n}{2^\ell}} \beta_{u,v} Z_u Z_v. \quad (121)$$

Also, due to the nature of the recursive decomposition, it suffices to focus on the implementation of a specific term such as $H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+2)\frac{n}{2^\ell}, (2b+3)\frac{n}{2^\ell}]}$. We will show momentarily that all the decomposed terms can be implemented with similar complexities $\text{cost}(\cdot)$, while the remaining terms such as $H_{[1+b\frac{n}{2^\ell}, (b+1)\frac{n}{2^\ell}]}$ and $H_{[1+b\frac{n}{2^\ell}, (b+1)\frac{n}{2^\ell}]:[1+(b+1)\frac{n}{2^\ell}, (b+2)\frac{n}{2^\ell}]}$ act on constant-size intervals and can be handled by a naive implementation using product formulas. This means the total complexity can be bounded as

$$\text{cost}_{\text{rec}}(n) = \sum_{\ell=2}^h 3 \left(2^{\ell-1} - 1 \right) \text{cost} \left(\frac{n}{2^{\ell-1}} \right) + \mathcal{O}(n), \quad (122)$$

which again reduces to the study of $\text{cost}(\cdot)$ because of the master theorem (Lemma 1). For simplicity, we choose $\ell = 2$, $b = 0$ and we study the complexity $\text{cost}(\frac{n}{2})$ of simulating $H_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}$.

Our main motivation to consider this decomposition is made clear through the following rank assumption of $H_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}$:

- (v) Coefficients of $H_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}$, when organized into an $\frac{n}{4}$ -by- $\frac{n}{4}$ matrix, have rank at most ρ , i.e., $\text{rank}([\beta_{u,v}]) \leq \rho$.

Because of this low-rank property, the coefficient matrix admits the thin singular value decomposition [56, Theorem 7.3.2]

$$\beta_{u,v} = \sum_{s=1}^{\rho} \mu_{u,s} \sigma_s \nu_{v,s} \quad (123)$$

for $u = 1 \dots, \frac{n}{4}$ and $v = \frac{n}{2} + 1, \dots, \frac{3n}{4}$, where μ and ν , when viewed as $\frac{n}{4}$ -by- ρ matrices, have real orthonormal columns, and σ_s are singular values bounded by the induced 1-norm of β and β^\top : $0 \leq \sigma_s \leq \sqrt{\|\beta\|_1 \|\beta^\top\|_1}$ [56, 5.6.P21]. Correspondingly, the exponential of $H_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}$ can be rewritten as

$$e^{-itH_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}} = e^{-it \sum_{s=1}^{\rho} \sigma_s \left(\sum_{u=1}^{\frac{n}{4}} \mu_{u,s} Z_u \right) \left(\sum_{v=\frac{n}{2}+1}^{\frac{3n}{4}} \nu_{v,s} Z_v \right)}. \quad (124)$$

With respect to the computational basis, our target exponential has the action

$$e^{-itH_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}} |z_n, \dots, z_1\rangle = e^{-it \sum_{s=1}^{\rho} \sigma_s \left(\sum_{u=1}^{\frac{n}{4}} \mu_{u,s} (-1)^{z_u} \right) \left(\sum_{v=\frac{n}{2}+1}^{\frac{3n}{4}} \nu_{v,s} (-1)^{z_v} \right)} |z_n, \dots, z_1\rangle. \quad (125)$$

This suggests a circuit implementation as follows. We first introduce an ancilla register of size $\mathcal{O}(\log(nt/\epsilon))$, and compute the following function in superposition

$$[z_n \ \cdots \ z_1] \mapsto \sum_{s=1}^{\rho} \sigma_s \left(\sum_{u=1}^{\frac{n}{4}} \mu_{u,s} (-1)^{z_u} \right) \left(\sum_{v=\frac{n}{2}+1}^{\frac{3n}{4}} \nu_{v,s} (-1)^{z_v} \right). \quad (126)$$

Note that the function values are real numbers of size at most $\mathcal{O}(\text{poly}(n))$ and can thus be faithfully stored in the $\mathcal{O}(\log(nt/\epsilon))$ -qubit ancilla register. We then use a sequence of controlled rotations to introduce the phase. We may now uncompute the ancilla register by reverting the circuit. This implements the desired exponential $e^{-itH_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}}$.

Our above circuit has a gate complexity of cost $\left(\frac{n}{2}\right) = \mathcal{O}(n\rho \text{polylog}(nt/\epsilon))$ for simulating one pair of intervals in the recursive decomposition. To implement the Trotter step, we redefine ρ to be the maximum rank of coefficient matrices, maximized over all pairs of intervals in the decomposition Eq. (57). Explicitly,

$$\rho := \max_{\sigma \neq \sigma' \in \{x, y, z\}} \max_{\ell=2, \dots, h} \max_{b=0, \dots, 2^{\ell-1}-2} \rho_{\ell, b}^{(\sigma, \sigma')}, \quad (127)$$

where $\rho_{\ell, b}$ is the largest truncation rank for the terms $H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+2)\frac{n}{2^\ell}, (2b+3)\frac{n}{2^\ell}]}$, $H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]}$, and $H_{[1+(2b+1)\frac{n}{2^\ell}, (2b+2)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]}$. By the master theorem (Lemma 1), this implies the circuit for Trotter steps has the same asymptotic cost

$$\text{cost}_{\text{rec}}(n) = \sum_{\ell=2}^h 3 \left(2^{\ell-1} - 1 \right) \text{cost} \left(\frac{n}{2^{\ell-1}} \right) + \mathcal{O}(n) = \mathcal{O} \left(n\rho \text{polylog} \left(\frac{nt}{\epsilon} \right) \right). \quad (128)$$

By simulating for time t/r in each step and repeating r steps where

$$r = \begin{cases} \frac{n^{o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}, & \alpha \geq 1, \\ \frac{n^{1-\alpha+o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}, & 0 < \alpha < 1, \end{cases} \quad (129)$$

we obtain the total gate complexity of the low-rank simulation method

$$\begin{cases} \frac{\rho(nt)^{1+o(1)}}{\epsilon^{o(1)}}, & \alpha \geq 1, \\ \frac{\rho n^{2-\alpha+o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}, & 0 < \alpha < 1. \end{cases} \quad (130)$$

Before ending this section, we discuss the important question of how the rank ρ is determined in the above decomposition. For various classes of power-law interactions including the Coulomb interaction, there are rigorous analyses based on the multipole expansion showing that $\rho = \mathcal{O}(\log(nt/\epsilon))$

suffices to guarantee that the simulation is ϵ -accurate,⁸ which leads to the gate complexity claimed in [Table 1](#). In fact, we will soon examine an application of this method to simulating the real-space electronic structure Hamiltonian in [Section 6](#). In practice, one may also run numerical simulations to empirically determine the rank value ρ . We refer the reader to [\[47\]](#) and papers citing this work for detailed studies of the low-rank decomposition in the classical setting.

5.2 Summary of the algorithm

We now summarize the low-rank method for implementing faster Trotter steps:

1. Use diagonalization to implement the low-rank matrix exponential [Eq. \(125\)](#).
2. Perform matrix exponentials for all combinations of Pauli operators $\sigma, \sigma' = x, y, z$, layers indexed by $\ell = 2, \dots, h$ and blocks indexed by $b = 0, \dots, 2^{\ell-1} - 2$, as well as the constant-size blocks, to implement a single Trotter step.
3. Repeat r Trotter steps to simulate the entire evolution.

Theorem 3 (Faster Trotter steps using low-rank decomposition). *Consider 2-local Hamiltonians $H = \sum_{\sigma \neq \sigma' \in \{x, y, z\}} \sum_{1 \leq j < k \leq n} \beta_{j,k}^{(\sigma, \sigma')} P_j^{(\sigma)} P_k^{(\sigma')}$, where $|\beta_{j,k}^{(\sigma, \sigma')}| \leq 1/|j - k|^\alpha$ for some constant $\alpha > 0$ and $P^{(\sigma)}$ ($\sigma = x, y, z$) are Pauli matrices. Let $t > 0$ be the simulation time and $\epsilon > 0$ be the target accuracy. Then H can be simulated with complexity*

$$\begin{cases} \rho n t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & \alpha \geq 1, \\ \rho n^{2-\alpha} t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & \alpha < 1. \end{cases} \quad (131)$$

Here, ρ defined in [Eq. \(127\)](#) is the maximum rank of certain off-diagonal blocks of coefficient matrices ($\rho = \mathcal{O}(\log(nt/\epsilon))$ if the coefficient distribution exactly matches a power law in one spatial dimension).

6 Applications to real-space quantum simulation

Simulating electronic structure Hamiltonians is one of the most widely studied problems in quantum simulation [\[24, 90, 93\]](#). An efficient solution of the electronic structure problem could lead to better understandings of catalysts and materials, which has applications in numerous subareas of physics and chemistry. Here, we consider mapping the electronic structure Hamiltonian on a grid and performing simulation in the second quantization in real space. We will present an algorithm combining our method for Trotter step implementation with a tighter error analysis, which improves the best simulation results from previous work.

We consider the following class of Hamiltonians

$$H = T + V := \sum_{j,k} \tau_{j,k} A_j^\dagger A_k + \sum_{l,m} \nu_{l,m} N_l N_m, \quad (132)$$

where A_j^\dagger, A_k are the fermionic creation and annihilation operators, N_l are the occupation-number operators, τ, ν are coefficient matrices, and the summations are over n spin orbitals. We can represent the real-space electronic structure Hamiltonians in the above form with specific choices

⁸Note that this logarithmic factor will be absorbed by $(nt/\epsilon)^{o(1)}$ in the final gate complexity.

of coefficients τ and ν ; we will come back to this point momentarily. Then, we simulate the Hamiltonian using product formulas by splitting e^{-itH} into products of e^{-itT} and e^{-itV} . The Trotter error corresponding to this splitting was studied in [111]. There, they found that a p th-order product formula $S_p(t)$ has error scaling like

$$\|S_p(t) - e^{-itH}\|_{\mathcal{W}_\eta} = \mathcal{O}\left((\|\tau\| + \|\nu\|_{\max} \eta)^{p-1} \|\tau\| \|\nu\|_{\max} \eta^2 t^{p+1}\right), \quad (133)$$

where recall $\|\cdot\|_{\max}$ is the max-norm denoting the largest matrix element in absolute value, $\|\cdot\|$ is the operator norm and

$$\|X\|_{\mathcal{W}_\eta} := \max_{|\psi_\eta\rangle, |\phi_\eta\rangle \in \mathcal{W}_\eta} |\langle \phi_\eta | X | \psi_\eta \rangle| \quad (134)$$

is the restriction of the operator norm to the η -electron subspace. Furthermore, if the coefficient matrices τ and ν are s -sparse (with at most s nonzero elements in each row and column), then it holds that

$$\|S_p(t) - e^{-itH}\|_{\mathcal{W}_\eta} = \mathcal{O}\left((\|\tau\|_{\max} + \|\nu\|_{\max})^{p-1} \|\tau\|_{\max} \|\nu\|_{\max} s^{p+1} \eta t^{p+1}\right). \quad (135)$$

The new bound we prove is as follows:

$$\|S_p(t) - e^{-itH}\|_{\mathcal{W}_\eta} = \mathcal{O}\left(\left(\|\tau\|_1 + \|\nu\|_{1, [\eta]}\right)^{p-1} \|\tau\|_1 \|\nu\|_{1, [\eta]} \eta t^{p+1}\right), \quad (136)$$

where we have used the induced 1-norm and its restricted version

$$\|\tau\|_1 := \max_j \sum_k |\tau_{j,k}|, \quad \|\nu\|_{1, [\eta]} := \max_j \max_{k_1 < \dots < k_\eta} (|\nu_{j,k_1}| + \dots + |\nu_{j,k_\eta}|). \quad (137)$$

This new bound is strictly better than Eq. (135) because for an s -sparse coefficient matrix β , it holds that $\|\beta\|_{1, [\eta]} \leq \|\beta\|_1 \leq s \|\beta\|_{\max}$ and the equalities are not always attainable. In fact, as we will see momentarily, the gap between the induced 1-norm $\|\beta\|_{1, [\eta]}$ and the max-norm $s \|\beta\|_{\max}$ is quite significant for the Coulomb interaction which is relevant for the electronic structure simulation. Compared with Eq. (133), our induced 1-norm scaling $\|\nu\|_{1, [\eta]}$ is again better than $\|\nu\|_{\max} \eta$, but our dependence on τ is slightly worse since $\|\tau\|_1 \geq \|\tau\|$ for a Hermitian τ . Fortunately, this is not an issue with the real-space electronic structure Hamiltonian. To avoid diluting the main message of our work, we state our Trotter error bound below and leave its proof to [Appendix B](#).

Theorem 4 (Trotter error with fermionic induced 1-norm scaling). *Let $H = T + V := \sum_{j,k} \tau_{j,k} A_j^\dagger A_k + \sum_{l,m} \nu_{l,m} N_l N_m$ be an interacting-electronic Hamiltonian, and $S_p(t)$ be a p th-order product formula splitting the evolutions under T and V . Then,*

$$\|S_p(t) - e^{-itH}\|_{\mathcal{W}_\eta} = \mathcal{O}\left(\left(\|\tau\|_1 + \|\nu\|_{1, [\eta]}\right)^{p-1} \|\tau\|_1 \|\nu\|_{1, [\eta]} \eta t^{p+1}\right). \quad (138)$$

Here, $\|\tau\|_1$ and $\|\nu\|_{1, [\eta]}$ are the (restricted) fermionic induced 1-norm defined by

$$\|\tau\|_1 = \max_j \sum_k |\tau_{j,k}|, \quad \|\nu\|_{1, [\eta]} = \max_j \max_{k_1 < \dots < k_\eta} (|\nu_{j,k_1}| + \dots + |\nu_{j,k_\eta}|). \quad (139)$$

We now apply the above theorem to the simulation of electronic structure Hamiltonians in real space. In this case, we have $H = T + U + V$, where T represents the kinetic term, U represents the external potential term introduced by the Born-Oppenheimer approximation, and V represents the

Coulomb potential term. We first consider the uniform electron gas without the external potential U . The Coulomb interaction takes the form

$$V = \frac{n^{1/3}}{2\omega^{1/3}} \sum_{l,m} \frac{1}{\|l-m\|} N_l N_m, \quad (140)$$

where n is the number of spin orbitals, ω is the volume of the computational cell, l and m are 3D vectors with each coordinate taking integer values between $-\frac{n^{1/3}-1}{2}$ and $\frac{n^{1/3}-1}{2}$. This term can be represented as $V = \sum_{l,m} \nu_{l,m} N_l N_m$ and it is easy to check that $\|\nu\|_{\max} = \mathcal{O}(n^{1/3}/\omega^{1/3})$. On the other hand, we can estimate the (restricted) induced 1-norm as follows. For a fixed value of l , the η spin orbitals nearest to l , if occupied, will have the largest possible coefficients $1/\|l-m\|$. These η spin orbitals form a ball of radius $\mathcal{O}(\eta^{1/3})$, which is also inside of a cube of linear size $\mathcal{O}(\eta^{1/3})$. This implies that [32, Lemma H.1]

$$\|\nu\|_{1, [\eta]} = \mathcal{O}\left(\frac{\eta^{2/3} n^{1/3}}{\omega^{1/3}}\right). \quad (141)$$

We thus see that the scaling of $\|\nu\|_{1, [\eta]}$ is strictly less than $\eta \|\nu\|_{\max}$, giving a factor of $\eta^{1/3}$ improvement over the best previous result. Meanwhile, the scaling for the kinetic term is⁹

$$\|\tau\|_1 = \mathcal{O}\left(\frac{n^{2/3}}{\omega^{2/3}}\right). \quad (142)$$

This determines the asymptotic scaling of the number of Trotter steps as

$$\left(\frac{\eta^{2/3} n^{1/3}}{\omega^{1/3}} + \frac{n^{2/3}}{\omega^{2/3}}\right) \frac{n^{o(1)} t^{1+o(1)}}{\epsilon^{o(1)}}. \quad (143)$$

To implement each Trotter step, one could naively exponentiate all terms in the Hamiltonian with cost $\Theta(n^2)$. However, note that under the Jordan-Wigner encoding, the Coulomb potential term is translated to the spin Hamiltonian

$$V = \frac{n^{1/3}}{2\omega^{1/3}} \sum_{l,m} \frac{1}{\|l-m\|} \frac{I - Z_l}{2} \frac{I - Z_m}{2}, \quad (144)$$

with interaction coefficients decaying with distance according to power law. This can be handled by the hierarchical low-rank method described in Section 5 for the 1D case with straightforward generalizations in Appendix C.1 to higher spatial dimensions. One can show using the multipole expansion that the decomposed terms can be approximated using matrices with rank polylogarithmic in the input parameters [47, 77]. So the complexity of implementing the Coulomb potential term is $\mathcal{O}(n \text{ polylog}(nt/\epsilon))$.

For the kinetic part, we can either represent it using a finite difference scheme [9], or using a diagonal form in the kinetic basis and transforming back [110]. In the first representation, we have geometrically local terms and we can further split them using product formulas without increasing the asymptotic scaling of Trotter error. Alternatively, we can modify boundary terms (using product formulas) so that the coefficients are cyclic, and the kinetic operator can then be

⁹The scaling we report here corresponds to the Fourier-based method introduced below. The finite-difference scheme of [9] seems to have a different scaling for the kinetic term, which should be further studied by future work.

implemented using the fermionic fast Fourier transform [9] (we can also use this implementation for the diagonal-form representation). As the first derivatives of the electronic wavefunction are not continuous at electron cusps [95], we note that the error introduced by even a first-order difference scheme should not differ asymptotically from high-order difference schemes, or even an exact diagonalization of the kinetic part, in the limit of large n . In either case, we can exponentiate the kinetic part with an asymptotic cost of $\mathcal{O}(n \text{ polylog}(nt/\epsilon))$ as well. This gives the total gate complexity

$$\left(\frac{\eta^{2/3} n^{4/3}}{\omega^{1/3}} + \frac{n^{5/3}}{\omega^{2/3}} \right) \frac{n^{o(1)} t^{1+o(1)}}{\epsilon^{o(1)}} \quad (145)$$

for the uniform electron gas, which simplifies to $\left(\eta^{1/3} n^{1/3} + \frac{n^{2/3}}{\eta^{2/3}} \right) n^{1+o(1)}$ assuming the density of electrons $\eta/\omega = \mathcal{O}(1)$ is fixed.

In second quantization, there were few studies about the asymptotic complexity of simulating the real-space electronic structure going beyond the Trotter error bound of [111] and the naive Trotter step implementation. So our result gives a significant improvement over the state-of-the-art scaling $\left(\eta^{2/3} n^{1/3} + \frac{n^{2/3}}{\eta^{2/3}} \right) n^{2+o(1)}$. In first quantization, recent work [110] gives an algorithm based on the interaction-picture simulation [87] with complexity $\tilde{\mathcal{O}}(\eta^{8/3} n^{1/3})$. So our result is asymptotically better as long as $n = \eta^{7/3 - o(1)}$.

For the general case, we have the external potential term

$$U = \sum_m \left(\sum_l \frac{\zeta_l}{\|\tilde{r}_l - r_m\|} \right) N_m \quad (146)$$

introduced under the Born-Oppenheimer approximation, where ζ_l are nuclear charges, \tilde{r}_l are nuclear coordinates, and r_m are the positions of electrons ($r_m = \frac{\omega^{1/3}}{n^{1/3}} m$). In this case, we will have a complexity similar to Eq. (145), but with an additional contribution from the external potential that depends on

$$\max_m \sum_l \frac{\zeta_l}{\|\tilde{r}_l - r_m\|}. \quad (147)$$

We note that this is strictly better than the interaction-picture algorithm from previous work [110] which instead depends on a larger quantity

$$\left(\sum_l \zeta_l \right) \max_{l,m} \frac{1}{\|\tilde{r}_l - r_m\|} = \mathcal{O} \left(\frac{\eta n^{1/3}}{\omega^{1/3}} \right). \quad (148)$$

A study of how much this improvement is for practical electronic structure problems will be left as a subject for future work.

7 First circuit lower bound

We have presented three methods for implementing faster Trotter steps and identified applications to electronic structure simulation in second quantization in real space. It is worth mentioning that all our methods make use of structural properties of the Hamiltonian in an essential way. Without such guarantees, the coefficient oracle in the block encoding and average-cost simulation described in Section 3 and Section 4 would have an implementation cost that depends on the number of terms $\Theta(n^2)$, whereas the blocks in the hierarchical decomposition of Section 5 would no longer have low

rank. Consequently, all our methods will reduce to the naive circuit implementation of Trotter steps.

In this section, we construct a class of 2-local Hamiltonians with only Pauli- Z interactions with coefficients taking a continuum range of values. We show that one needs at least $\Omega(n^2)$ gates to evolve this class of Hamiltonians with accuracy $\epsilon = \Omega(1/\text{poly}(n))$ for time $t = \Omega(\epsilon)$. This suggests that using structural properties of the Hamiltonian is both necessary and sufficient for performing faster Trotter steps. We prove a related circuit lower bound in [Appendix A](#) to justify the necessity of structural properties to efficiently implement the coefficient oracle O_β with approximate circuit synthesis.

The proof of our circuit lower bound is based on a gate-efficient reduction to the problem of approximately synthesizing diagonal unitaries in the Hamming weight-2 subspace, which we will describe in detail in [Section 7.2](#). With this strategy in mind, we first examine the approximate synthesis of diagonal unitaries in [Section 7.1](#).

7.1 Approximate synthesis of diagonal unitaries

We now consider the gate complexity lower bound for synthesizing diagonal unitaries. The exact version of this problem was considered by Bullock and Markov [\[21\]](#). Here, we generalize their lower bound to the approximate synthesis, by adapting a volume comparison technique of Knill [\[67\]](#).

To be specific, our goal is to implement μ -qubit diagonal unitaries with accuracy δ

$$|x_\mu, \dots, x_2, x_1\rangle \mapsto e^{i\theta_x} |x_\mu, \dots, x_2, x_1\rangle, \quad (149)$$

where the phase angles take values up to $|\theta_x| \leq \theta_{\max} < \pi$. Recall that the set of all such unitaries is denoted by $\mathcal{D}_{\theta_{\max}}$. Our circuit acts on a total number of $b \geq \mu$ qubits. We will first prove a lower bound for a 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$, and then bootstrap the result to obtain a lower bound for arbitrary 2-qubit gates.

As aforementioned, this lower bound will be proved using a volume comparison argument: we will compare the volume of all the diagonal unitaries within δ -distance to the set of constructable quantum circuits, against the volume of our target set of diagonal unitaries. A similar argument was used by Knill to prove a circuit lower bound for synthesizing the full unitary group $\mathcal{U}(2^\mu)$ [\[67\]](#), but our proof here is significantly simplified as we restrict to the set of diagonal unitaries whose volume is easier to compute (see [Section 2.4](#)).

We start by bounding the number of b -qubit circuits that can be constructed using g gates from a finite two-qubit gate set \mathcal{K} . For each of the g gates, we have $\binom{b}{2}$ choices for the locations in which the two-qubit gate acts and, fixing the locations, we have $|\mathcal{K}|$ choices for the specific gate. This give the upper bound

$$\left(\binom{b}{2} |\mathcal{K}| \right)^g \quad (150)$$

for the number of distinct circuits we can construct.

Now, to synthesize diagonal unitaries with accuracy δ , we need all of them to fall into the δ -ball of some quantum circuits (in spectral norm distance). For each circuit V , the volume of diagonal unitaries within its δ -ball can be estimated as follows. We take any two diagonal unitaries U and \tilde{U} from the δ -ball centered at V . By the triangle inequality,

$$\|U - V\|, \|\tilde{U} - V\| \leq \delta \quad \Rightarrow \quad \|\tilde{U} - U\| \leq 2\delta, \quad (151)$$

which implies

$$\left| e^{i\tilde{\theta}_x} - e^{i\theta_x} \right| \leq 2\delta \quad (152)$$

when restricting to the specific diagonal elements indexed by x . Solving for $\tilde{\theta}_x$ and θ_x when $\delta < 1$, we have

$$\left| \tilde{\theta}_x - \theta_x \right| \leq \arcsin \left(2\delta \sqrt{1 - \delta^2} \right). \quad (153)$$

This means $\max_x \max_{\tilde{\theta}_x, \theta_x} \left| \tilde{\theta}_x - \theta_x \right| \leq \arcsin \left(2\delta \sqrt{1 - \delta^2} \right)$, so the volume of such diagonal unitaries is upper bounded by $\left(\arcsin \left(2\delta \sqrt{1 - \delta^2} \right) \right)^{2\mu}$. We thus have that the total volume of synthesizable diagonal unitaries is at most

$$\begin{aligned} & \text{vol}(\{U|U \text{ is diagonal unitary, } \exists \text{ circuit } V, \|U - V\| \leq \delta\}) \\ & \leq \left(\binom{b}{2} |\mathcal{K}| \right)^g \max_{\text{circuit } V} \text{vol}(\{U|U \text{ is diagonal unitary, } \|U - V\| \leq \delta\}) \\ & \leq \left(\binom{b}{2} |\mathcal{K}| \right)^g \text{vol} \left(\mathcal{D}_{\frac{\arcsin(2\delta\sqrt{1-\delta^2})}{2}} \right) = \left(\binom{b}{2} |\mathcal{K}| \right)^g \left(\arcsin \left(2\delta \sqrt{1 - \delta^2} \right) \right)^{2\mu}. \end{aligned} \quad (154)$$

On the other hand, we know that the total volume of the target diagonal unitaries is

$$\text{vol}(\mathcal{D}_{\theta_{\max}}) = (2\theta_{\max})^{2\mu}. \quad (155)$$

For our synthesis to succeed, we require that

$$\{U|U \text{ is diagonal unitary, } \exists \text{ circuit } V, \|U - V\| \leq \delta\} \supseteq \mathcal{D}_{\theta_{\max}}, \quad (156)$$

which implies

$$\left(\binom{b}{2} |\mathcal{K}| \right)^g \left(\arcsin \left(2\delta \sqrt{1 - \delta^2} \right) \right)^{2\mu} \geq (2\theta_{\max})^{2\mu}, \quad g \geq 2\mu \cdot \frac{\log \left(\frac{2\theta_{\max}}{\arcsin(2\delta\sqrt{1-\delta^2})} \right)}{\log \left(\binom{b}{2} |\mathcal{K}| \right)}. \quad (157)$$

Note that this lower bound is only useful when $\theta_{\max} = \Omega(\delta)$, for otherwise the right-hand side of the bound becomes negative. Intuitively, we would simply use the identity operator in the circuit synthesis if the target diagonal unitary itself is close to identity. On the other hand, the restriction $\theta_{\max} < \pi$ is not fundamental: if $\theta_{\max} \geq \pi$, we can simply reset it to a smaller value.

We have so far proved a circuit lower bound when the underlying gate set \mathcal{K} is of finite size. If we are allowed to use arbitrary 2-qubit gates, then the bound can be modified as follows. We first synthesize the circuit with respect to a fixed finite universal gate set \mathcal{K}' , say Clifford+T. This gives a new circuit approximating the original one to accuracy δ with a slightly larger gate complexity $cg \log(g/\delta)$ for some constant $c > 0$. We then invoke the above bound with the number of gates $cg \log(g/\delta)$, accuracy 2δ and the gate set $\mathcal{K}' = \text{Clifford+T}$ of constant size. This gives

$$cg \log \left(\frac{g}{\delta} \right) = \Omega \left(2\mu \cdot \frac{\log \left(\frac{\theta_{\max}}{\delta} \right)}{\log b} \right). \quad (158)$$

Solving g using the Lambert-W function [44, Lemma 59], we obtain the same asymptotic scaling

$$g = \Omega \left(2\mu \cdot \frac{\log \left(\frac{\theta_{\max}}{\delta} \right)}{\log b} \right). \quad (159)$$

We summarize this bound in the following theorem.

Theorem 5 (Approximate synthesis of diagonal unitaries). *Consider diagonal unitaries $\mathcal{D}_{\theta_{\max}} = \{\sum_x e^{i\theta_x} |x_\mu, \dots, x_1\rangle\langle x_\mu, \dots, x_1| \mid |\theta_x| \leq \theta_{\max} < \pi\}$ on μ qubits with phase angles at most $\theta_{\max} < \pi$. Given accuracy $0 < \delta < 1$, number of qubits $b \geq \mu$, and 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$,*

$$\begin{aligned} & \min \{g \mid \forall U \in \mathcal{D}_{\max}, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from set } \mathcal{K}, \|U - V\| \leq \delta\} \\ & \geq 2^\mu \cdot \frac{\log\left(\frac{2\theta_{\max}}{\arcsin(2\delta\sqrt{1-\delta^2})}\right)}{\log\left(\binom{b}{2}|\mathcal{K}|\right)}. \end{aligned} \quad (160)$$

Under the same assumption (with $0 < \delta < 1/2$ and $\theta_{\max} = \Omega(\delta)$) but choosing \mathcal{K} to be the set of arbitrary 2-qubit gates,

$$\begin{aligned} & \min \{g \mid \forall U \in \mathcal{D}_{\max}, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from set } \mathcal{K}, \|U - V\| \leq \delta\} \\ & = \Omega\left(\frac{2^\mu}{\log b}\right). \end{aligned} \quad (161)$$

7.2 Reduction from simulation in the Hamming weight-2 subspace

We now show a circuit lower bound for simulating a class of 2-local Hamiltonians with only Pauli- Z interactions. Specifically, we consider $H = \sum_{1 \leq j < k \leq n} \beta_{j,k} Z_j Z_k$, where coefficients take a continuum range of values up to $|\beta_{j,k}| \leq t$. Suppose first that this evolution can be simulated to accuracy δ using circuits acting on $b \geq n$ qubits with gates chosen from a 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$. An analogous lower bound holds when the gate set contains arbitrary 2-qubit gates.

The key technical ingredient of our proof is a gate-efficient reduction from the approximate synthesis of diagonal unitaries within the Hamming weight-2 subspace \mathcal{W}_2 . To elaborate, we rewrite the evolution as

$$e^{-iH} = e^{i \sum_{j < k} \beta_{j,k} I} e^{-i \sum_{j < k} \beta_{j,k} Z_j} e^{-i \sum_{j < k} \beta_{j,k} Z_k} e^{-i \sum_{j < k} 4\beta_{j,k} \frac{I - Z_j}{2} \frac{I - Z_k}{2}}. \quad (162)$$

It suffices for us to consider only the last exponential since all the remaining ones can be performed with gate complexity $\mathcal{O}(n)$ (with only logarithmic overhead when compiled with respect to a fixed finite universal gate set, say Clifford+T). We have

$$\sum_{j < k} \beta_{j,k} \frac{I - Z_j}{2} \frac{I - Z_k}{2} |e_u \oplus e_v\rangle = \beta_{u,v} |e_u \oplus e_v\rangle, \quad (163)$$

for all $1 \leq u < v \leq n$, where e_u represents the n -bit string with 1 at the u th position and 0 elsewhere and \oplus denotes the modulo-2 addition. Using a conversion mapping that will be introduced below, we can then relate the problem of simulating 2-local Hamiltonians with coefficients $\beta_{j,k}$, to the synthesis of diagonal unitaries with phase angles $\beta_{j,k}$, within the Hamming weight-2 subspace.

To complete the reduction, we need to convert the usual computational basis states to the basis states of the Hamming weight-2 subspace \mathcal{W}_2 . This can be done using an inequality test, followed by a binary-to-unary conversion. Specifically, suppose we have two quantum registers each of length $\log n$ holding the binary representation of j and k (assuming n is a power of 2). We then use an inequality test with cost $\mathcal{O}(\log n)$ to implement

$$|j\rangle|k\rangle \mapsto |j\rangle|k\rangle|j?k\rangle = \begin{cases} |j\rangle|k\rangle|j = k\rangle, & j = k, \\ |j\rangle|k\rangle|j < k\rangle, & j < k, \\ |j\rangle|k\rangle|j > k\rangle, & j > k. \end{cases} \quad (164)$$




Figure 4: Illustration of the reduction. This circuit implements diagonal unitaries on registers $|j\rangle$ and $|k\rangle$ by simulating 2-local commuting Hamiltonians on an ancilla register (shaded orange). The reduction is gate efficient, so the asymptotic cost of quantum simulation is the same as that of synthesizing diagonal unitaries.

We can introduce arbitrary phases for the first case using $\mathcal{O}(n)$ gates (which can again be compiled using Clifford+T with only logarithmic overhead). For the case where $j < k$, we apply the binary-to-unary conversion with cost $\mathcal{O}(n)$ to get

$$|j\rangle|k\rangle|j < k\rangle \mapsto |j\rangle|k\rangle|e_j \oplus e_k\rangle|j < k\rangle, \quad (165)$$

after which we introduce the desired phases by simulating the corresponding 2-local Hamiltonian with accuracy ϵ . The last case $j > k$ can be handled similarly by swapping the role of j and k . See Figure 4 for an illustration of this reduction.

To summarize, we can use the above circuit to implement arbitrary diagonal unitaries on the two quantum registers $|j\rangle|k\rangle$ of total length $\log n + \log n = \log(n^2)$ to accuracy 3ϵ with phase angles taking values up to $|\beta_{j,k}| \leq t$. Our circuit uses two instances of controlled simulations of 2-local Hamiltonians, together with $\mathcal{O}(n \text{ polylog}(n/\epsilon))$ additional Clifford+T gates, acting on a total number of $\mathcal{O}(b)$ qubits (recall $b \geq n$). In other words, if the simulation uses g gates from \mathcal{K} , then we can implement arbitrary diagonal unitaries using

$$2g + \mathcal{O}(n \text{ polylog}(n/\epsilon)) \quad (166)$$

gates from the gate set

$$\mathcal{K}' = \text{controlled-}\mathcal{K} \cup (\text{Clifford} + \text{T}). \quad (167)$$

However, our above lower bound indicates that

$$\Omega\left(\frac{n^2}{\log(b|\mathcal{K}'|)}\right) \quad (168)$$

gates are necessary to implement diagonal unitaries acting on $b = \log(n^2)$ qubits with accuracy 3ϵ for $t = \Omega(\epsilon)$. Thus, similar asymptotic bound applies to the uncontrolled simulation as well by observing that $|\mathcal{K}'| = \Theta(|\mathcal{K}|)$.

Corollary 1 (Simulating 2-local commuting Hamiltonians). *Consider 2-local Hamiltonians $H = \sum_{1 \leq j < k \leq n} \beta_{j,k} Z_j Z_k$, where coefficients take values up to $|\beta_{j,k}| \leq t$. Given accuracy $0 < \epsilon < 1/3$, number of qubits $b \geq n$, and 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$, if $t = \Omega(\epsilon)$,*

$$\begin{aligned} & \min \{g \mid \forall \text{ 2-local Hamiltonian } H, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from } \mathcal{K}, \|e^{-iH} - V\| \leq \epsilon\} \\ & = \Omega\left(\frac{n^2}{\log(b|\mathcal{K}|)} - n \text{ polylog}\left(\frac{n}{\epsilon}\right)\right). \end{aligned} \quad (169)$$

Under the same assumption but choosing \mathcal{K} to be the set of arbitrary 2-qubit gates,

$$\begin{aligned} & \min \{g \mid \forall \text{ 2-local Hamiltonian } H, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from } \mathcal{K}, \|e^{-iH} - V\| \leq \epsilon\} \\ & = \Omega \left(\frac{n^2}{\log b} - n \text{ polylog} \left(\frac{n}{\epsilon} \right) \right). \end{aligned} \tag{170}$$

To apply our bound for implementing Trotter steps, we first divide the evolution into r steps and simulate each step using product formulas for time t/r with accuracy ϵ/r . Depending on the specific definition of the Hamiltonian, we may perform further decompositions like in [Section 3](#), [Section 4](#) and [Section 5](#) and apply the bound only to commuting terms. We note that when the Hamiltonian coefficients have nonuniform magnitudes, the minimum effective evolution time can often be much smaller. For instance, the effective time for power-law interactions scale like $\Theta\left(\frac{t}{n^{\alpha r}}\right)$ as the smallest terms have size $\Theta(1/n^\alpha)$. To apply our bound, we would then require that $t = \Omega(n^\alpha \epsilon)$, which is a pessimistic estimate. Of course, one can replace the volume formula in [Eq. \(155\)](#) with a better estimate that takes account of the nonuniformness of the Hamiltonian coefficients. However, this modification will not lead to a significant improvement to our bound as the dependence on the magnitude of coefficients is only logarithmic. We leave a detailed study of such improvement for future work.

8 Discussion

In this work, we have developed circuit implementations of Trotter steps going beyond the naive approach that cycles through all terms in the Hamiltonian. We have presented three methods, one based on the block encoding technique, one based on an average-cost simulation and one based on a recursive low-rank decomposition. This allows us to simulate systems with power-law decaying interactions with significantly lower cost. As an application, we have given an algorithm for simulating electronic structure Hamiltonians in real space, with complexity $\left(\eta^{1/3}n^{1/3} + \frac{n^{2/3}}{\eta^{2/3}}\right)n^{1+o(1)}$ for the uniform electron gas and a similar cost when the external potential is included, improving over the best results from previous work. In achieving these speedups, we have made essential use of structural properties of the system. We further construct a class of 2-local commuting Hamiltonians on n qubits with continuum range of coefficients whose simulation requires $\Omega(n^2)$ gates. Our results thus suggest that the use of structural properties of the Hamiltonian is both necessary and sufficient to implement Trotter steps with lower asymptotic gate complexity.

Although we have employed different techniques to develop various faster simulation algorithms, the core idea behind all our improvements is the use of recursion. With suitable recursive decompositions, we have showed that the target Hamiltonian can be significantly simplified with lower 1-norm and smaller rank value, both of which are useful for quantum simulation. Meanwhile, because of the master theorem ([Lemma 1](#)), recursions only introduce logarithmic factors to the overall gate complexity, thus simplifying the target problem ‘‘almost for free’’. We expect that recursions will find further applications in reducing the cost of simulating many-body Hamiltonians.

We have focused on simulating 2-local spin models throughout our paper, and it would be interesting to study general κ -local models as well as fermionic systems. For this purpose, previous work [[94](#)] utilized different low-rank decompositions for 1- and 2-body fermionic operators which led to Trotter steps with lower cost. However, an issue not rigorously addressed in their work is that the decomposition does not respect the commutation relations of Hamiltonian terms and could lead to a Trotter error larger than the naive approach. Thus, some additional efforts are needed to justify the utility of those methods in quantum simulation. Note that this drawback can be

slightly alleviated using the phase gradient construction [64] if the fault-tolerant resource cost is of interest. On the other hand, our techniques do apply to other κ -local models for constant κ and to fermionic systems, as long as the Hamiltonian can be recursively decomposed as in Section 3, Section 4 and Section 5. Preliminary studies of such generalizations were available in the context of measuring quantum observables [18], and we hope future work could develop more efficient Hamiltonian decompositions along this line.

We have demonstrated that advanced quantum simulation techniques, such as block encoding and qubitization, can be used to improve the performance of product formulas. Such techniques can be directly applied to implement Trotter steps; however, this would introduce a slow-down factor proportional to the 1-norm of Hamiltonian coefficients, offering no benefit over the naive circuit implementation. We have showed how to overcome this obstacle for power-law interactions by recursively decomposing the Hamiltonian using product formulas to reduce the norm and by implementing an average-case cost simulation. Our result is complementary to recent results on the hybridized simulation [101, 112], where quantum algorithms are “hybridized” with the interaction-picture method to achieve better performance. Our finding indicates that product formulas can be just as effective in hybridizing with other algorithms for quantum simulation.

In our low-rank method, we have performed a recursive decomposition of the target Hamiltonian where certain off-diagonal blocks of coefficients have rank value logarithmic in the input parameters. Such a decomposition was used in a different context by [96] to block-encode kernel matrices. Although their results are potentially useful for improving electronic structure simulation in first quantization, we are not aware of a simple circuit to realize this idea. In our problem, the decomposed Hamiltonian terms can be handled directly, so block encoding is no longer needed (it would be inefficient anyway due to the normalization factor issue mentioned above). We apply this method in Section 6 to implement the Coulomb potential term of the electronic structure Hamiltonian in second quantization in real space, obtaining improvements over the best previous simulation results.

When no structural property is available in the target Hamiltonian, then intuitively the best method one can hope for is to exponentiate the terms one by one. We justify this intuition in Section 7 with a circuit lower bound, showing that a class of 2-local commuting Hamiltonians on n sites requires at least $\Omega(n^2)$ gates to simulate. However, we emphasize that our lower bound is useful only when the effective evolution time is sufficiently long: if the effective time is short compared to the accuracy threshold, our bound becomes trivial. However, such failures are often indications that improved implementations of Trotter steps exist, as one can truncate the Hamiltonian and replace small-angle exponentials by identity operators. Previous work has studied various truncation schemes for product formulas as well as more advanced quantum algorithms [14, 32, 33, 58, 91, 118]. By exploiting the relative scaling between evolution time and system size, it is plausible to get faster quantum simulation algorithms with better dependence on the spacetime volume.

Although our lower bound is mostly used to argue the difficulty of implementing product formulas, it is applicable to other simulation algorithms as well. In particular, it applies to the sampling-based approach [23], as one can always reproduce classical computations on a quantum computer, and the use of randomness can be represented unitarily via the Stinespring dilation. A similar argument then lower bounds both the classical and quantum resources used by the simulation. But given the difficulty of realizing a scalable quantum computer, it would be beneficial if certain computations can be offloaded classically, which is exactly a main goal of the sampling-based algorithm. Looking forward, it would be interesting to further study the tradeoffs between classical and quantum resources in designing quantum algorithms.

Several other related questions are worth further investigation. To simplify our analysis, we have used boxes with exponentially growing sizes to group each qubit index in the block encoding

approach, but this may be improved by using different box definitions. We have also chosen a uniform division of intervals in the realization of average-cost simulation, but improvements might be possible through a nonuniform division scheme. We have almost exclusively considered the number of arbitrary 2-qubit gates as the cost metric for quantum simulation, and future work may consider other metrics such as circuit depth, locality [97] and fault-tolerant gate count (although we expect some of our techniques are useful under those metrics as well). With these questions in mind, we hope our work will motivate continuing progress on finding the most efficient methods to solve instances of quantum simulation problems with practical interest.

Acknowledgements

Y.S. thanks Ryan Babbush for helpful discussions during the initial stages of this work. Y.T. is supported in part by the U.S. Department of Energy Office of Science (DE-SC0019374), Office of Advanced Scientific Computing Research (DE-SC0020290), Office of High Energy Physics (DE-ACO2-07CH11359), and under the Quantum System Accelerator project. M.C.T. is supported by the Defense Advanced Research Project Agency (DARPA) under contract number 134371-5113608 and the Quantum Algorithms and Machine Learning grant from Nippon Telegraph and Telephone (NTT), number AGMT DTD 9/24/20.

A Second circuit lower bound

In [Section 7](#), we proved a lower bound on the gate complexity of approximately synthesizing diagonal unitaries, and bootstrapped it to get a lower bound on the simulation of 2-local commuting Hamiltonians. In this appendix, we prove a related lower bound on the cost of synthesizing the coefficient oracle O_β used in [Section 3](#) and [Section 4](#) for block encoding. To this end, we first prove a lower bound for the approximate synthesis of discrete diagonal unitaries in [Appendix A.1](#), and then do a reduction from applications of the coefficient oracle in [Appendix A.2](#). Our result extends a previous bound of [85] by allowing approximate circuit synthesis and use of arbitrary 2-qubit gates.

A.1 Approximate synthesis of discrete diagonal unitaries

We first consider the gate complexity of synthesizing discrete diagonal unitaries of the form

$$|x\rangle \mapsto e^{i\frac{\beta_{x,m-1}2^{m-1} + \dots + \beta_{x,0}}{2^m}} |x\rangle, \quad (171)$$

with accuracy δ . Here, x takes 2^μ possible values, and β_x is an m -bit string parameterized by x (so $\beta_{x,j} \in \{0,1\}$). We use a setup similar to that of [Section 7](#). Specifically, our circuit acts on a total number of $b \geq \mu$ qubits. We will first prove a lower bound for a 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$, and then bootstrap the result to obtain a lower bound for arbitrary 2-qubit gates.

As before, the number of distinct b -qubit circuits that can be constructed using g gates from gate set \mathcal{K} is at most

$$\left(\binom{b}{2} |\mathcal{K}| \right)^g. \quad (172)$$

We now repeat our volume argument. To synthesize all discrete diagonal unitaries with accuracy δ , we need all of them to fall into the spectral norm δ -ball of some quantum circuits. Fixing a specific

circuit V , we consider any diagonal unitaries U and \tilde{U} from the δ -ball centered at V . Because of the triangle inequality, we must have

$$\|\tilde{U} - U\| \leq 2\delta. \quad (173)$$

However, the new observation is that our target set is discrete and finite. For $\delta = \mathcal{O}(1)$ and $\delta = \Omega\left(\frac{1}{2^m}\right)$, these unitaries have the x th diagonal entries

$$\left| e^{i\frac{\tilde{\beta}_{x,m-1}2^{m-1}+\dots+\tilde{\beta}_{x,0}}{2^m}} - e^{i\frac{\beta_{x,m-1}2^{m-1}+\dots+\beta_{x,0}}{2^m}} \right| \leq 2\delta. \quad (174)$$

Solving for the (decimal) $\tilde{\beta}_x$ and β_x , we have

$$\left| \tilde{\beta}_x - \beta_x \right| \leq 2^m \arcsin\left(2\delta\sqrt{1-\delta^2}\right), \quad (175)$$

which implies that the number of such discrete diagonal unitaries is at most

$$\mathcal{O}\left((2^m\delta)^{2^\mu}\right). \quad (176)$$

Using the counting measure $\#\{\cdot\}$, we thus have that the total number of synthesizable discrete diagonal unitaries is at most

$$\begin{aligned} & \#\{U|U \text{ is discrete diagonal unitary}, \exists \text{ circuit } V, \|U - V\| \leq \delta\} \\ & \leq \left(\binom{b}{2}|\mathcal{K}|\right)^g \max_{\text{circuit } V} \#\{U|U \text{ is discrete diagonal unitary}, \|U - V\| \leq \delta\} \\ & = \left(\binom{b}{2}|\mathcal{K}|\right)^g \mathcal{O}\left((2^m\delta)^{2^\mu}\right) = \mathcal{O}\left(b^{2g}|\mathcal{K}|^g 2^{m2^\mu} \delta^{2^\mu}\right). \end{aligned} \quad (177)$$

On the other hand, by a counting argument, we know that the total number of discrete diagonal unitaries is exactly

$$\#\{U|U \text{ is discrete diagonal unitary}\} = (2^m)^{2^\mu} = 2^{m2^\mu}. \quad (178)$$

For our synthesis to succeed, we require that

$$\{U \text{ is discrete diagonal unitary} \mid \exists \text{ circuit } V, \|U - V\| \leq \delta\} \supseteq \{U \text{ is discrete diagonal unitary}\}, \quad (179)$$

which implies

$$g = \Omega\left(2^\mu \frac{\log\left(\frac{1}{\delta}\right)}{\log(b|\mathcal{K}|)}\right). \quad (180)$$

If \mathcal{K} is the arbitrary 2-qubit gate set, we can use the Lambert-W function in a similar way as in [Section 7.1](#) to get

$$g = \Omega\left(2^\mu \frac{\log\left(\frac{1}{\delta}\right)}{\log b}\right). \quad (181)$$

We have thus proved:

Theorem 6 (Approximate synthesis of discrete diagonal unitaries). *Consider discrete diagonal unitaries $\left\{ \sum_x e^{i \frac{\beta_{x,m-1} 2^{m-1} + \dots + \beta_{x,0}}{2^m}} |x\rangle\langle x| \right\}$ on μ qubits with phase angles determined by m -bit strings β_x . Given accuracy $\delta = \mathcal{O}(1)$ and $\delta = \Omega\left(\frac{1}{2^m}\right)$, number of qubits $b \geq \mu$, and 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$,*

$$\begin{aligned} & \min \{g \mid \forall \text{discrete diagonal unitary } U, \exists b\text{-qubit circuit } V \text{ with } g \text{ gates from set } \mathcal{K}, \|U - V\| \leq \delta\} \\ &= \Omega\left(\frac{2^\mu \log\left(\frac{1}{\delta}\right)}{\log(b|\mathcal{K}|)}\right). \end{aligned} \tag{182}$$

Under the same assumption but choosing \mathcal{K} to be the set of arbitrary 2-qubit gates,

$$\begin{aligned} & \min \{g \mid \forall \text{discrete diagonal unitary } U, \exists b\text{-qubit circuit } V \text{ with } g \text{ gates from set } \mathcal{K}, \|U - V\| \leq \delta\} \\ &= \Omega\left(\frac{2^\mu \log\left(\frac{1}{\delta}\right)}{\log b}\right). \end{aligned} \tag{183}$$

A.2 Reduction from the coefficient oracle

Recall that in the block encoding method, we have used the following oracle of Hamiltonian coefficients:

$$O_\beta |u, v, 0\rangle = |u, v, \beta_{u,v}\rangle. \tag{184}$$

Here, u and v each takes n values and $\beta_{u,v}$ are m -bit binary strings. To ensure the efficiency of our method, we have assumed that the gate complexity of implementing oracle O_β is $\mathcal{O}(\text{polylog}(nt/\epsilon))$. Implicitly, this implies the existence of certain structural properties of $\beta_{u,v}$: we now show that if $\beta_{u,v}$ are allowed to depend arbitrarily on u and v , then one needs at least $\sim n^2$ gates to implement this oracle, resulting in an inefficient block encoding approach.

The key ingredient of the proof is a reduction from the application of coefficient oracle to the approximate synthesis of discrete diagonal unitaries. This is done as follows:

$$\begin{aligned} |u, v, 0\rangle & \xrightarrow{O_\beta} |u, v, \beta_{u,v}\rangle \\ & \mapsto e^{i \frac{\tilde{\beta}_{u,v}}{2^m}} |u, v, \beta_{u,v}\rangle \\ & \xrightarrow{O_\beta^\dagger} e^{i \frac{\tilde{\beta}_{u,v}}{2^m}} |u, v, 0\rangle, \end{aligned} \tag{185}$$

where in the second step we introduce the approximate phase factors $\tilde{\beta}_{u,v}$ by applying $\Theta(\log(1/\epsilon))$ rotations $|0\rangle\langle 0| + e^{i\frac{\tilde{\beta}_{u,v}}{2^m}} |1\rangle\langle 1|$ ($\ell = 1, \dots, \Theta(\log(1/\epsilon))$). We have thus implemented an arbitrary discrete diagonal unitary on registers $|u, v\rangle$ of total length $\mu = 2 \log n$ using a circuit of 2 oracle calls (each with accuracy ϵ) and $\Theta(\log(1/\epsilon))$ single-qubit gates (which can be compiled with respect to a fixed gate set like Clifford+T with only logarithmic overhead).

In summary, suppose O_β can be implemented to accuracy ϵ using a b -qubit circuit of g gates chosen from a gate set \mathcal{K} of finite size $|\mathcal{K}|$. Then, we can implement the entire reduction to accuracy 3ϵ using a b -qubit circuit of $2g + \mathcal{O}(\text{polylog}(1/\epsilon))$ gates chosen from $\mathcal{K}' = \mathcal{K} \cup (\text{Clifford} + \text{T})$. But our lower bound indicates that

$$\Omega\left(\frac{n^2 \log\left(\frac{1}{\epsilon}\right)}{\log(b|\mathcal{K}'|)}\right) = \Omega\left(\frac{n^2 \log\left(\frac{1}{\epsilon}\right)}{\log(b|\mathcal{K}|)}\right) \tag{186}$$

gates are required for the synthesis task. We have thus proved:

Corollary 2 (Approximate synthesis of coefficient oracles). *Consider the set of coefficient oracles $\{O_\beta \mid O_\beta|u, v, 0\rangle = |u, v, \beta_{u,v}\rangle\}$ where u and v each takes n values and $\beta_{u,v}$ are m -bit strings. Given accuracy $\epsilon = \mathcal{O}(1)$ and $\epsilon = \Omega\left(\frac{1}{2^m}\right)$, number of qubits $b \geq m + 2\log n$, and 2-qubit gate set \mathcal{K} of finite size $|\mathcal{K}|$,*

$$\begin{aligned} & \min \{g \mid \forall \text{oracle } O_\beta, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from set } \mathcal{K}, \|O_\beta - V\| \leq \epsilon\} \\ &= \Omega\left(\frac{n^2 \log\left(\frac{1}{\epsilon}\right)}{\log(b|\mathcal{K}|)} - \text{polylog}\left(\frac{1}{\epsilon}\right)\right). \end{aligned} \quad (187)$$

Under the same assumption but choosing \mathcal{K} to be the set of arbitrary 2-qubit gates,

$$\begin{aligned} & \min \{g \mid \forall \text{oracle } O_\beta, \exists \text{ circuit } V \text{ on } b \text{ qubits with } g \text{ gates from set } \mathcal{K}, \|O_\beta - V\| \leq \epsilon\} \\ &= \Omega\left(\frac{n^2 \log\left(\frac{1}{\epsilon}\right)}{\log b} - \text{polylog}\left(\frac{1}{\epsilon}\right)\right). \end{aligned} \quad (188)$$

We see that implementations of the coefficient oracle O_β will generally cost $\sim n^2$ gates as well, even when approximate circuit synthesis is allowed. Thus, one needs structural properties of the Hamiltonian not only to realize faster Trotter steps but to implement faster coefficient oracles as well.

B Trotter error with fermionic induced 1-norm scaling

In this appendix, we prove the fermionic induced 1-norm scaling of Trotter error claimed in [Theorem 4](#), which is in turn used in the simulation of real-space electronic structure Hamiltonians. Specifically, we consider the Hamiltonian

$$H = T + V := \sum_{j,k} \tau_{j,k} A_j^\dagger A_k + \sum_{l,m} \nu_{l,m} N_l N_m, \quad (189)$$

where A_j^\dagger, A_k are the fermionic creation and annihilation operators, N_l are the occupation-number operators, τ, ν are coefficient matrices, and the summations are over n spin orbitals. Let $S_p(t)$ be a p th-order product formula that splits e^{-itH} into products of e^{-itT} and e^{-itV} . Our goal is to bound the Trotter error $\|S_p(t) - e^{-itH}\|_{\mathcal{W}_\eta}$ restricted to the η -electron subspace \mathcal{W}_η (spanned by basis states with Hamming weight η).

To describe our proof, we introduce some additional notations and terminologies. For $\gamma = 0, 1$, we let

$$H^{(\gamma)} = \sum_{j,k} \mu_{j,k}^{(\gamma)} H_{j,k}^{(\gamma)}, \quad \mu^{(0)} = \nu, \quad \mu^{(1)} = \tau, \quad H_{j,k}^{(0)} = N_j N_k, \quad H_{j,k}^{(1)} = A_j^\dagger A_k. \quad (190)$$

In other words, $H^{(0)} = V, H^{(1)} = T$. Then, we have the commutator scaling of Trotter error from [Lemma 2](#):

$$\|S_p(t) - e^{-itH}\|_{\mathcal{W}_\eta} = \mathcal{O}\left(\max_{\gamma \in \{0,1\}^{p+1}} \max_{|\psi_\eta\rangle \in \mathcal{W}_\eta} \left| \langle \psi_\eta | \left[H^{(\gamma_{p+1})}, \dots \left[H^{(\gamma_2)}, H^{(\gamma_1)} \right] \right] | \psi_\eta \rangle \right| t^{p+1} \right), \quad (191)$$

where we have taken the expectation value since nested commutators of Hermitian operators are necessarily anti-Hermitian. Note the following commutation rules for fermionic operators [54]

$$\left[A_j^\dagger A_k, A_{j_x}^\dagger \right] = \delta_{k,j_x} A_j^\dagger, \quad \left[A_j^\dagger A_k, A_{k_y} \right] = -\delta_{k_y,j} A_k, \quad \left[A_j^\dagger A_k, N_{l_z} \right] = \delta_{k,l_z} A_j^\dagger A_{l_z} - \delta_{j,l_z} A_{l_z}^\dagger A_k. \quad (192)$$

$$\left[N_l N_m, A_{j_x}^\dagger \right] = \delta_{m,j_x} N_l A_{j_x}^\dagger + \delta_{l,j_x} A_{j_x}^\dagger N_m, \quad \left[N_l N_m, A_{k_x} \right] = -\delta_{m,k_x} N_l A_{k_x} - \delta_{l,k_x} A_{k_x} N_m, \quad (193)$$

where $\delta_{j,k} = 1$ if and only if $j = k$. Thus, the nested commutator $\left[H_{j_{p+1},k_{p+1}}^{(\gamma_{p+1})}, \dots, \left[H_{j_2,k_2}^{(\gamma_2)}, H_{j_1,k_1}^{(\gamma_1)} \right] \right]$ can be written as a linear combination of products of elementary fermionic operators. Following [111, Proposition 11], we call each summand P from the nested commutator a *fermionic path* and write $P \triangleright \left(H_{j_{p+1},k_{p+1}}^{(\gamma_{p+1})}, \dots, H_{j_1,k_1}^{(\gamma_1)} \right)$. Then, we have the following bound

$$\begin{aligned} & \left| \langle \psi_\eta | \left[H^{(\gamma_{p+1})}, \dots, \left[H^{(\gamma_2)}, H^{(\gamma_1)} \right] \right] | \psi_\eta \rangle \right| \\ & \leq \max_{\|c\|_1 = \eta} \sum_{j_{p+1},k_{p+1}} \dots \sum_{j_1,k_1} \sum_{P \triangleright \left(H_{j_{p+1},k_{p+1}}^{(\gamma_{p+1})}, \dots, H_{j_1,k_1}^{(\gamma_1)} \right)} \left| \mu_{j_{p+1},k_{p+1}}^{(\gamma_{p+1})} \dots \mu_{j_1,k_1}^{(\gamma_1)} \right| \frac{\|P|c\rangle\| + \|P^\dagger|c\rangle\|}{2}, \end{aligned} \quad (194)$$

where the maximization is taken over all n -bit strings c with Hamming weight η , and $P|c\rangle$ ($P^\dagger|c\rangle$) either gives another basis state or gives the zero vector.

Previous work [111] further proceeds to upper bound the coefficients by the max-norm $\|\tau\|_{\max}$ and $\|\nu\|_{\max}$ and counts the number of fermionic paths with nonzero contribution to the summation. This results in a bound with the max-norm scaling asymptotically worse than our result. To get our improved **Theorem 4**, we need a more careful treatment of the summation order as well as the action of fermionic operators on a fixed basis state.

We now present a proof by induction. The base case of $p = 0$ is easy to check. For any n -bit string c with Hamming weight η , we have

$$\begin{aligned} & \sum_{j,k} |\tau_{j,k}| \frac{\|A_j^\dagger A_k |c\rangle\| + \|A_k^\dagger A_j |c\rangle\|}{2} = \mathcal{O}(\|\tau\|_1 \eta), \\ & \sum_{l,m} |\nu_{l,m}| \frac{\|N_l N_m |c\rangle\| + \|N_m N_l |c\rangle\|}{2} = \mathcal{O}(\|\nu\|_{1,[\eta]} \eta). \end{aligned} \quad (195)$$

This is because in order to make a nonzero contribution, the right-most fermionic operators only have η possible choices. This motivates us to choose the summation ordering so that the indices of operators on the left are summed over first:

$$\begin{aligned} & \sum_k \sum_j |\tau_{j,k}| \|A_j^\dagger A_k |c\rangle\|, & \sum_j \sum_k |\tau_{j,k}| \|A_k^\dagger A_j |c\rangle\|, \\ & \sum_m \sum_l |\nu_{l,m}| \|N_l N_m |c\rangle\|, & \sum_l \sum_m |\nu_{l,m}| \|N_m N_l |c\rangle\|. \end{aligned} \quad (196)$$

Fixing one choice for the right-most operator, the remaining quantity can be upper bounded by either the induced 1-norm $\|\tau\|_1$ or the restricted induced 1-norm $\|\nu\|_{1,[\eta]}$ respectively. This proves the desired bound in the base case.

For the inductive step, assume that given any fermionic path $P \triangleright \left(H_{j_p,k_p}^{(\gamma_p)}, \dots, H_{j_1,k_1}^{(\gamma_1)} \right)$, we can always reorder the summation to yield the desired bound. Now take the next layer of commutator

$[T, P]$ or $[V, P]$. By Eq. (192) and Eq. (193) and a separate induction, $P = \prod_{q=1}^{p+2} P_q$ is the product of at most $p + 2$ elementary fermionic operators (which include A_j^\dagger, A_k, N_l), and at any point in the product the number of creation operators is always smaller than the number of annihilation operators. Then, the commutation may be performed sequentially as

$$[T, P] = \sum_{q=1}^{p+2} P_{p+2} \cdots P_{q+1} [T, P_q] P_{q-1} \cdots P_1, \quad [V, P] = \sum_{q=1}^{p+2} P_{p+2} \cdots P_{q+1} [V, P_q] P_{q-1} \cdots P_1. \quad (197)$$

Depending on $P_q = A_{j_x}^\dagger, A_{k_y}$ or N_{l_z} , we divide the analysis into 5 subcases:

- (i) $[A_j^\dagger A_k, A_{j_x}^\dagger] = \delta_{k, j_x} A_j^\dagger$: We keep the original sum over j_x while introducing a new sum over j **on the right**. The summation range of j is set to be over **all n spin orbitals**. The number of electrons increases by 1 due to A_j^\dagger with the new location determined by j (as opposed to j_x). Asymptotically, this does not affect the summation of other existing indices and increases the norm by a factor of $\|\tau\|_1$ through the index j .
- (ii) $[A_j^\dagger A_k, A_{k_y}] = -\delta_{k_y, j} A_k$: We keep the original sum over k_y while introducing a new sum over k **on the right**. The summation range of k is set to be over **all n spin orbitals**. The number of electrons decreases by 1 due to A_k with the new location determined by k (as opposed to k_y). Asymptotically, this does not affect the summation of other existing indices and increases the norm by a factor of $\|\tau\|_1$ through the index k .
- (iii) $[A_j^\dagger A_k, N_{l_z}] = \delta_{k, l_z} A_j^\dagger A_{l_z} - \delta_{j, l_z} A_{l_z}^\dagger A_k$: We keep the original sum over l_z while introducing a new sum over j (k) **on the right** respectively. The summation range of j (k) is set to be over **all n spin orbitals**. The number of electrons remains the same but the occupation changes with the locations determined by j (k) and l_z . Asymptotically, this does not affect the summation of other existing indices and increases the norm by a factor of $\|\tau\|_1$ through the index j (k).
- (iv) $[N_l N_m, A_{j_x}^\dagger] = \delta_{m, j_x} N_l A_{j_x}^\dagger + \delta_{l, j_x} A_{j_x}^\dagger N_m = \delta_{m, j_x} N_l A_{j_x}^\dagger + \delta_{l, j_x} N_m A_{j_x}^\dagger - \delta_{l, j_x} \delta_{m, j_x} A_{j_x}^\dagger$: We keep the original sum over j_x while introducing a new sum over l (m) **on the right** respectively. The summation range of l (m) is set to be over **all occupied spin orbitals**. Both the number and the locations of electrons remain the same. Asymptotically, this does not affect the summation of other existing indices and increases the norm by a factor of $\|\nu\|_{1, [\eta]}$ through the index l (m).
- (v) $[N_l N_m, A_{k_x}] = -\delta_{m, k_x} N_l A_{k_x} - \delta_{l, k_x} A_{k_x} N_m = -\delta_{m, k_x} N_l A_{k_x} - \delta_{l, k_x} N_m A_{k_x} - \delta_{l, k_x} \delta_{m, k_x} A_{k_x}$: We keep the original sum over k_x while introducing a new sum over l (m) **on the right** respectively. The summation range of l (m) is set to be over **all occupied spin orbitals**. Both the number and the locations of electrons remain the same. Asymptotically, this does not affect the summation of other existing indices and increases the norm by a factor of $\|\nu\|_{1, [\eta]}$ through the index l (m).

Note that in the fourth and fifth cases, we have implicitly used the fact that the number of occupied modes is no larger than η . The proof of [Theorem 4](#) is now complete.

C Generalization

C.1 Higher spatial dimensionality

In the main body of the paper, we have presented faster methods to perform Trotter steps for simulating power-law interactions in one spatial dimension. Here, we briefly describe how to generalize those methods to higher spatial dimensionality.

Specifically, consider an n -qubit d -dimensional square lattice $\mathcal{L} \subseteq \mathbb{Z}^d$. Without loss of generality, we assume $n^{1/d}$ is a power of 2 and take $\mathcal{L} = \{1, \dots, n^{1/d}\}^d$. Then, the Hamiltonian has the form $H = \sum_{j,k \in \mathcal{L}} H_{j,k}$. For power-law interactions, we have the norm condition $\|H_{j,k}\| \leq 1/\|j-k\|^\alpha$ for constant $\alpha > 0$, where $\|j-k\|$ is simply the Euclidean distance between d -dimensional vectors j and k . Just like in the 1D case, we can expand the Hamiltonian with respect to tensor products of Pauli operators and hereafter assume the Hamiltonian takes the form

$$H = \sum_{\substack{j,k \in \mathcal{L} \\ j \neq k}} \beta_{j,k} X_j Y_k. \quad (198)$$

The commutator norm Eq. (28) corresponding to such decompositions has the scaling [32, Theorem H.2]

$$\tilde{\alpha}_{\text{comm}}^{o(1)} = \begin{cases} n^{o(1)}, & \alpha \geq d, \\ n^{1-\frac{\alpha}{d}+o(1)}, & 0 < \alpha < d, \end{cases} \quad (199)$$

which implies

$$r = \begin{cases} \frac{n^{o(1)} t^{1+o(1)}}{e^{o(1)}}, & \alpha \geq d, \\ \frac{n^{1-\frac{\alpha}{d}+o(1)} t^{1+o(1)}}{e^{o(1)}}, & 0 < \alpha < d. \end{cases} \quad (200)$$

Depending on the order $j_s < k_s$ or $j_s > k_s$ between each pair of the corresponding coordinates, we can further divide the Hamiltonian into 2^d subterms. The remaining degenerate cases where some pair of indices collide $j_s = k_s$ can be merged accordingly. We assume that

$$H = \sum_{1 \leq j_1 < k_1 \leq n^{1/d}} \cdots \sum_{1 \leq j_d < k_d \leq n^{1/d}} \beta_{j,k} X_j Y_k \quad (201)$$

in the following discussion, and a similar analysis holds for other cases with some notation changes.

We now perform a recursive decomposition along each spatial dimension of the system. So for the first spatial dimension, we have the term

$$\sum_{\substack{1 \leq j_1 < \frac{n^{1/d}}{2} \\ \frac{n^{1/d}}{2} \leq k_1 \leq n^{1/d}}} \cdots \sum_{1 \leq j_d < k_d \leq n^{1/d}} \beta_{j,k} X_j Y_k \quad (202)$$

in the first layer and terms

$$\sum_{\substack{1 \leq j_1 < \frac{n^{1/d}}{4} \\ \frac{n^{1/d}}{4} \leq k_1 < \frac{n^{1/d}}{2}}} \cdots \sum_{1 \leq j_d < k_d \leq n^{1/d}} \beta_{j,k} X_j Y_k, \quad \sum_{\substack{\frac{n^{1/d}}{2} \leq j_1 < \frac{3n^{1/d}}{4} \\ \frac{3n^{1/d}}{4} \leq k_1 \leq n^{1/d}}} \cdots \sum_{1 \leq j_d < k_d \leq n^{1/d}} \beta_{j,k} X_j Y_k \quad (203)$$

in the second layer and so on. Because of the master theorem [Lemma 1](#), we may restrict to a single term in the decomposition, say the one in the first layer. We then perform another recursive




Figure 5: Illustration of the 2D recursive decomposition. Decomposed terms act across two disjoint cubes. The colors indicate types of the operator support following the convention in Figure 1. Subfigure (a) represents terms in the first two layers of the decomposition of the first spatial dimension. Subfigure (b) recursively decomposes the first layer from (a) in the second spatial dimension.




Figure 6: Illustration of an average-cost block encoding for power-law Hamiltonians in 2D.

decomposition along the second spatial dimension, and once again restrict to only a specific term in the decomposition using Lemma 1. We iterate through all the d spatial dimensions. See Figure 5 for an illustration of this strategy in the 2D case.

For the block encoding method in Section 3, the cost of circuit implementation depends on the 1-norm of Hamiltonian coefficients, which determines the number of steps in the qubitization algorithm. Specifically, the cost of implementing decomposed Hamiltonian terms should now be revised to

$$\text{cost}(n) = \begin{cases} \mathcal{O}\left(\left(\frac{t}{r} + 1\right) (\sqrt{\lambda_{\text{block}}} + n) \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & \alpha \geq 2d, \\ \mathcal{O}\left(\left(n^{2-\frac{\alpha}{d}} \frac{t}{r} + 1\right) (\sqrt{\lambda_{\text{block}}} + n) \text{polylog}\left(\frac{nt}{\epsilon}\right)\right), & 0 < \alpha < 2d, \end{cases} \quad (204)$$

which implies that the entire simulation has a complexity of

$$\begin{cases} (\sqrt{\lambda_{\text{block}}} + n) t \left(\frac{nt}{\epsilon}\right)^{o(1)}, & \alpha \geq 2d, \\ (\sqrt{\lambda_{\text{block}}} + n) \left(n^{2-\frac{\alpha}{d}} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + t \left(\frac{nt}{\epsilon}\right)^{o(1)}\right), & d \leq \alpha < 2d, \\ (\sqrt{\lambda_{\text{block}}} + n) \left(n^{2-\frac{\alpha}{d}} t \text{polylog}\left(\frac{nt}{\epsilon}\right) + n^{1-\frac{\alpha}{d}} t \left(\frac{nt}{\epsilon}\right)^{o(1)}\right), & 0 < \alpha < d. \end{cases} \quad (205)$$

This simulation can be further improved for $\alpha < 2d$ using the average-cost simulation technique described in Section 4. Specifically, for the two cubes sharing a common vertex, we further divide each spatial dimension into m intervals of equal length; see Figure 6 for an illustration of this decomposition in 2D. Then, considering evolution time t/r , we have the modified cost function

$$\text{cost}(n) = \tilde{\mathcal{O}}\left(\left(n^{2-\frac{\alpha}{d}} \frac{t}{r} + m^{2d}\right) \left(\sqrt{\lambda_{\text{avg}}} + \frac{n}{m^d}\right)\right), \quad (206)$$

which implies that the entire simulation has cost

$$\left\{ \begin{array}{l} \min \left\{ \sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{d}} t \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} t + n^{2-\frac{\alpha}{2d}} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)}, \right. \\ \quad \left. \left(\sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{d}} t + n^{3-\frac{\alpha}{d}} t \right) \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} t + nt \right) \left(\frac{nt}{\epsilon} \right)^{o(1)} \right\}, \\ \min \left\{ \sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{d}} t \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{1-\frac{\alpha}{d}} t + n^{\frac{5}{2}-\frac{\alpha}{d}} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)}, \right. \\ \quad \left. \left(\sqrt{\lambda_{\text{avg}}} n^{2-\frac{\alpha}{d}} t + n^{3-\frac{\alpha}{d}} t \right) \text{polylog} \left(\frac{nt}{\epsilon} \right) + \left(\sqrt{\lambda_{\text{avg}}} n^{1-\frac{\alpha}{d}} t + n^{2-\frac{\alpha}{d}} t \right) \left(\frac{nt}{\epsilon} \right)^{o(1)} \right\}, \end{array} \right. \quad \begin{array}{l} d \leq \alpha < 2d, \\ 0 < \alpha < d. \end{array} \quad (207)$$

The generalization of the low-rank method in [Section 5](#) is similar and straightforward, by using higher-dimensional cubes instead of 1D intervals to perform the recursive decomposition. Unlike the block encoding method, here we do not have the issue of normalization factor; the circuit implementation is just a direct translation of classical computations. Moreover, for Coulomb interactions in higher spatial dimensions, one can show using the multipole expansion that the decomposed terms can also be approximated with a matrix with rank logarithmic in the input parameters. The resulting implementation has a complexity similar to that of the 1D case:

$$\begin{cases} \rho nt \left(\frac{nt}{\epsilon} \right)^{o(1)}, & \alpha \geq d, \\ \rho n^{2-\frac{\alpha}{d}} t \left(\frac{nt}{\epsilon} \right)^{o(1)}, & \alpha < d. \end{cases} \quad (208)$$

We refer the reader to [\[47\]](#) for a more detailed discussion of the decomposition.

C.2 κ -local Hamiltonians

We have focused on the simulation of 2-local Hamiltonians in the main body of the paper. In this subsection, we discuss how to extend our approach to handle more general κ -local Hamiltonians for constant κ .

Specifically, the Hamiltonian takes the form $H = \sum_{1 \leq j_1 < j_2 < \dots < j_\kappa \leq n} H_{j_1, j_2, \dots, j_\kappa}$, where operator $H_{j_1, j_2, \dots, j_\kappa}$ acts nontrivially only on qubits j_1, \dots, j_κ . Similar to [Eq. \(48\)](#), we can rewrite the Hamiltonian using tensor products of Pauli operators and decompose the evolution accordingly using product formulas. This introduces a Trotter error no larger than the naive decomposition where all Hamiltonian terms are decomposed. Thus without loss of generality, we consider

$$H = \sum_{1 \leq j_1 < j_2 < \dots < j_\kappa \leq n} \beta_{j_1, j_2, \dots, j_\kappa} P_{j_1}^{(\sigma_1)} P_{j_2}^{(\sigma_2)} \dots P_{j_\kappa}^{(\sigma_\kappa)}, \quad (209)$$

where $P_{j_b}^{(\sigma_b)}$ is a Pauli operator on qubit j labeled by $\sigma_b = x, y, z$.

To describe the recursive decomposition, we introduce the abbreviation

$$\begin{aligned} H_{[j, k]} &:= \sum_{j \leq u_1 < \dots < u_\kappa \leq k} \beta_{u_1, \dots, u_\kappa} P_{u_1}^{(\sigma_1)} \dots P_{u_\kappa}^{(\sigma_\kappa)}, \\ H_{[j_1, k_1]: [j_2, k_2]: \dots: [j_\kappa, k_\kappa], \kappa_1: \kappa_2: \dots: \kappa_\kappa} &:= \sum_{j_1 \leq u_1 < \dots < u_{\kappa_1} \leq k_1} \sum_{j_2 \leq v_1 < \dots < v_{\kappa_2} \leq k_2} \dots \sum_{j_\kappa \leq w_1 < \dots < w_{\kappa_\kappa} \leq k_\kappa} \\ &\quad \beta_{u_1, \dots, u_{\kappa_1}, v_1, \dots, v_{\kappa_2}, \dots, w_1, \dots, w_{\kappa_\kappa}} P_{u_1}^{(\sigma_1)} \dots P_{w_{\kappa_\kappa}}^{(\sigma_\kappa)} \end{aligned} \quad (210)$$

to represent terms within a specific interval and across κ disjoint intervals of sites, where $j_1 \leq k_1 < j_2 \leq k_2 < \dots$ and $\kappa_1, \dots, \kappa_\kappa$ are nonnegative integers summing to $\kappa_1 + \dots + \kappa_\kappa = \kappa$. Unlike the 2-local case, we now need additional parameters $\kappa_1, \dots, \kappa_\kappa$ to record the number of indices within each interval. With that, the decomposition used in the block encoding method in [Section 3](#) should be modified to

$$\begin{aligned}
H_{[1,n]} = & \sum_{\substack{\kappa_1 + \dots + \kappa_\kappa = \kappa \\ \forall \kappa_b < \kappa}} H_{[1, \frac{n}{\kappa}] : [\frac{n}{\kappa} + 1, 2\frac{n}{\kappa}] : \dots : [(\kappa-1)\frac{n}{\kappa}, n], \kappa_1 : \kappa_2 : \dots : \kappa_\kappa} \\
& + \sum_{b=0}^{\kappa-1} H_{[b\frac{n}{\kappa} + 1, (b+1)\frac{n}{\kappa}]} .
\end{aligned} \tag{211}$$

We have assumed that κ is also a power of 2, but the general case can be handled by redefining the boundary terms. Now, terms in the second line can be unwrapped under the same recursion, whereas terms in the first line correspond to degenerate cases with locality parameter $< \kappa$ (there are at most $\binom{2\kappa-1}{\kappa-1} = \mathcal{O}(1)$ cases) that can be handled by an outer induction on κ . We may modify the decompositions in [Section 4](#) and [Section 5](#) in a similar way. However, the Hamiltonian coefficients now have a more complicated tensor structure, so further studies are required to understand the complexity of implementing Trotter steps using such methods.

C.3 General fermionic Hamiltonians

Fermionic Hamiltonians that are geometrically local can be mapped to spin models using a locality-preserving encoding such as [\[114\]](#) and solved using methods from the previous subsection. We now discuss how to potentially apply our techniques to simulating general fermionic Hamiltonians. For presentational purpose, we only examine the one-body operator in one spatial dimension:

$$H = \sum_{j,k=1}^n \beta_{j,k} A_j^\dagger A_k, \tag{212}$$

where A_j^\dagger and A_k are creation and annihilation operators on fermionic modes j and k , respectively. The analysis of higher spatial dimensionality and many-body fermionic operators may proceed in a similar way as in [Appendix C.1](#) and [Appendix C.2](#).

We introduce the following abbreviation

$$\begin{aligned}
H_{[j,k]} & := \sum_{j \leq u < v \leq k} \beta_{u,v} \left(A_u^\dagger A_v + A_v^\dagger A_u \right) \quad (1 \leq j < k \leq n), \\
H_{[j,k]:[l,m]} & := \sum_{\substack{j \leq u \leq k \\ l \leq v \leq m}} \beta_{u,v} \left(A_u^\dagger A_v + A_v^\dagger A_u \right) \quad (1 \leq j \leq k < l \leq m \leq n).
\end{aligned} \tag{213}$$

This is similar to that used in [Section 3](#), [Section 4](#) and [Section 5](#), except we have added the complex-conjugate terms to make the entire operator Hermitian. For the recursive decomposition with a reduced 1-norm, we have

$$H = \sum_{\ell=1}^{\log n - 1} \sum_{b=0}^{2^{\ell-1} - 1} H_{[2b\frac{n}{2^\ell} + 1, (2b+1)\frac{n}{2^\ell}] : [(2b+1)\frac{n}{2^\ell} + 1, 2(b+1)\frac{n}{2^\ell}]} . \tag{214}$$

This decomposition shares similar features as that in [Section 3](#). In particular, the 1-norm of each summand is significantly smaller compared to the full Hamiltonian, suggesting that the block encoding method can be advantageous.

We now show how to realize an efficient block encoding. Without loss of generality, we choose $\ell = 1$, $b = 0$ and study the term $H_{[1, \frac{n}{2}]:[\frac{n}{2}+1, n]}$. This can be block-encoded as follows. We define the Majorana operators

$$M_{u,0} := A_u^\dagger + A_u, \quad M_{u,1} := i \left(A_u^\dagger - A_u \right), \quad (215)$$

using which we rewrite

$$\begin{aligned} H_{[1, \frac{n}{2}]:[\frac{n}{2}+1, n]} &= \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} \beta_{u,v} \left(A_u^\dagger A_v + A_v^\dagger A_u \right) \\ &= \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} \beta_{u,v} \left(\frac{M_{u,0} - iM_{u,1}}{2} \frac{M_{v,0} + iM_{v,1}}{2} + \frac{M_{v,0} - iM_{v,1}}{2} \frac{M_{u,0} + iM_{u,1}}{2} \right) \\ &= \sum_{\substack{1 \leq u \leq \frac{n}{2} \\ \frac{n}{2}+1 \leq v \leq n}} i\beta_{u,v} (M_{u,0}M_{v,1} + M_{v,0}M_{u,1}). \end{aligned} \quad (216)$$

The preparation subroutine can now be implemented in a similar way as in [Section 3.2](#). For the selection subroutine, we need to iterate over products of Majorana operators, which can be realized with cost $\mathcal{O}(n)$ (see for example [\[8, Section III.B\]](#)).

For the recursive decomposition described in [Section 5](#), we have

$$\begin{aligned} H &= \sum_{\ell=2}^{\eta} \sum_{b=0}^{2^{\ell-1}-2} \left(H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+2)\frac{n}{2^\ell}, (2b+3)\frac{n}{2^\ell}]} \right. \\ &\quad + H_{[1+2b\frac{n}{2^\ell}, (2b+1)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]} \\ &\quad \left. + H_{[1+(2b+1)\frac{n}{2^\ell}, (2b+2)\frac{n}{2^\ell}]:[1+(2b+3)\frac{n}{2^\ell}, (2b+4)\frac{n}{2^\ell}]} \right) \\ &\quad + \sum_{b=0}^{2^\eta-1} H_{[1+b\frac{n}{2^\ell}, (b+1)\frac{n}{2^\ell}]} + \sum_{b=0}^{2^\eta-2} H_{[1+b\frac{n}{2^\ell}, (b+1)\frac{n}{2^\ell}]:[1+(b+1)\frac{n}{2^\ell}, (b+2)\frac{n}{2^\ell}]} \end{aligned} \quad (217)$$

Without loss of generality, we choose $\ell = 2$, $b = 0$ and study the term

$$H_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]} = \sum_{\substack{1 \leq u \leq \frac{n}{4} \\ \frac{n}{2}+1 \leq v \leq \frac{3n}{4}}} \beta_{u,v} \left(A_u^\dagger A_v + A_v^\dagger A_u \right) \quad (218)$$

with the coefficient matrix

$$\tilde{\beta} = \begin{bmatrix} 0 & 0 & \beta & 0 \\ 0 & 0 & 0 & 0 \\ \beta^\top & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (219)$$

Note that although the Hamiltonian terms no long commute in the fermionic case, if $\text{rank}(\beta) = \rho$, then $\text{rank}(\tilde{\beta}) = \rho$ as well, and we can still implement the exponential of $H_{[1, \frac{n}{4}]:[\frac{n}{2}+1, \frac{3n}{4}]}$ using fermionic Givens rotations with cost $\mathcal{O}(\rho n)$ [\[56, 2.1.P28\]](#) (see [\[65\]](#) for the circuit implementation). Of course, the circuit complexity will depend on the target Hamiltonian as well as the decomposition scheme, which become more complicated to analyze for general fermionic operators. We leave a detailed study of these generalizations as a subject for future work.

References

- [1] Nilin Abrahamsen, Yuan Su, Yu Tong, and Nathan Wiebe, *Entanglement area law for 1D gauge theories and bosonic systems*, 2022, [arXiv:2203.16012](#).
- [2] Dong An, Di Fang, and Lin Lin, *Time-dependent unbounded Hamiltonian simulation with vector norm scaling*, *Quantum* **5** (2021), 459, [arXiv:2012.13105](#).
- [3] Dong An and Lin Lin, *Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm*, *ACM Transactions on Quantum Computing* **3** (2022), no. 2, 1–28, [arXiv:1909.05500](#).
- [4] Abhinav Anand, Philipp Schleich, Sumner Alperin-Lea, Phillip W. K. Jensen, Sukin Sim, Manuel Daz-Tinoco, Jakob S. Kottmann, Matthias Degroote, Artur F. Izmaylov, and Alán Aspuru-Guzik, *A quantum computing view on unitary coupled cluster theory*, *Chemical Society Reviews* **51** (2022), 1659–1684, [arXiv:2109.15176](#).
- [5] Anurag Anshu, Aram W. Harrow, and Mehdi Soleimanifar, *Entanglement spread area law in gapped ground states*, *Nature Physics* **18** (2022), no. 11, 1362–1366, [arXiv:2004.15009](#).
- [6] Yosi Atia and Dorit Aharonov, *Fast-forwarding of Hamiltonians and exponentially precise measurements*, *Nature Communications* **8** (2017), no. 1, 1572, [arXiv:1610.09619](#).
- [7] Ryan Babbush, Dominic W. Berry, Jarrod R. McClean, and Hartmut Neven, *Quantum simulation of chemistry with sublinear scaling in basis size*, *npj Quantum Information* **5** (2019), no. 1, 92, [arXiv:1807.09802](#).
- [8] Ryan Babbush, Dominic W. Berry, and Hartmut Neven, *Quantum simulation of the Sachdev-Ye-Kitaev model by asymmetric qubitization*, *Physical Review A* **99** (2019), 040301, [arXiv:1806.02793](#).
- [9] Ryan Babbush, Nathan Wiebe, Jarrod McClean, James McClain, Hartmut Neven, and Garnet Kin-Lic Chan, *Low-depth quantum simulation of materials*, *Physical Review X* **8** (2018), 011044, [arXiv:1706.00023](#).
- [10] Christian W. Bauer, Zohreh Davoudi, A. Baha Balantekin, Tanmoy Bhattacharya, Marcela Carena, Wibe A. de Jong, Patrick Draper, Aida El-Khadra, Nate Gemelke, Masanori Hanada, Dmitri Kharzeev, Henry Lamm, Ying-Ying Li, Junyu Liu, Mikhail Lukin, Yannick Meurice, Christopher Monroe, Benjamin Nachman, Guido Pagano, John Preskill, Enrico Rinaldi, Alessandro Roggero, David I. Santiago, Martin J. Savage, Irfan Siddiqi, George Siopsis, David Van Zanten, Nathan Wiebe, Yukari Yamauchi, Kbra Yeter-Aydeniz, and Silvia Zorzetti, *Quantum simulation for high energy physics*, 2022, [arXiv:2204.03381](#).
- [11] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders, *Efficient quantum algorithms for simulating sparse Hamiltonians*, *Communications in Mathematical Physics* **270** (2007), no. 2, 359–371, [arXiv:quant-ph/0508139](#).
- [12] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma, *Exponential improvement in precision for simulating sparse Hamiltonians*, *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pp. 283–292, 2014, [arXiv:1312.1414](#).

- [13] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma, *Simulating Hamiltonian dynamics with a truncated Taylor series*, Physical Review Letters **114** (2015), no. 9, 090502, [arXiv:1412.4687](#).
- [14] 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** (2019), 208, [arXiv:1902.02134](#).
- [15] Dominic W Berry, Maria Kieferová, Artur Scherer, Yuval R Sanders, Guang Hao Low, Nathan Wiebe, Craig Gidney, and Ryan Babbush, *Improved techniques for preparing eigenstates of fermionic Hamiltonians*, npj Quantum Information **4** (2018), no. 1, 22, [arXiv:1711.10460](#).
- [16] Sergio Blanes and Fernando Casas, *A concise introduction to geometric numerical integration*, Chapman and Hall/CRC, 2016.
- [17] Alex Bocharov, Martin Roetteler, and Krysta M. Svore, *Efficient synthesis of universal Repeat-Until-Success quantum circuits*, Physical Review Letters **114** (2015), 080502, [arXiv:1404.5320](#).
- [18] Xavier Bonet-Monroig, Ryan Babbush, and Thomas E. O’Brien, *Nearly optimal measurement scheduling for partial tomography of quantum states*, Physical Review X **10** (2020), 031064, [arXiv:1908.05628](#).
- [19] Joseph W. Britton, Brian C. Sawyer, Adam C. Keith, C. C Joseph Wang, James K. Freericks, Hermann Uys, Michael J. Biercuk, and John J. Bollinger, *Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins*, Nature **484** (2012), 489, [arXiv:1204.5789](#).
- [20] H. P. Büchler, E. Demler, M. Lukin, A. Micheli, N. Prokof’ev, G. Pupillo, and P. Zoller, *Strongly correlated 2D quantum phases with cold polar molecules: Controlling the shape of the interaction potential*, Physical Review Letters **98** (2007), 060404, [arXiv:cond-mat/0607294](#).
- [21] Stephen S. Bullock and Igor L. Markov, *Asymptotically optimal circuits for arbitrary n -qubit diagonal computations*, Quantum Information & Computation **4** (2004), no. 1, 2747, [arXiv:quant-ph/0303039](#).
- [22] Daniel Burgarth, Niklas Galke, Alexander Hahn, and Lauritz van Luijk, *State-dependent Trotter limits and their approximations*, 2022, [arXiv:2209.14787](#).
- [23] Earl Campbell, *Random compiler for fast Hamiltonian simulation*, Physical Review Letters **123** (2019), 070503, [arXiv:1811.08017](#).
- [24] Yudong Cao, Jonathan Romero, Jonathan P. Olson, Matthias Degroote, Peter D. Johnson, Mária Kieferová, Ian D. Kivlichan, Tim Menke, Borja Peropadre, Nicolas P. D. Sawaya, Sukin Sim, Libor Veis, and Alán Aspuru-Guzik, *Quantum chemistry in the age of quantum computing*, Chemical Reviews **119** (2019), no. 19, 10856–10915, [arXiv:1812.09976](#).
- [25] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery, *The power of block-encoded matrix powers: Improved regression techniques via faster Hamiltonian simulation*, 46th International Colloquium on Automata, Languages, and Programming, vol. 132, pp. 33:1–33:14, 2019, [arXiv:1804.01973](#).

- [26] Hans Hon Sang Chan, Richard Meister, Tyson Jones, David P. Tew, and Simon C. Benjamin, *Grid-based methods for chemistry simulations on a quantum computer*, 2022, [arXiv:2202.05864](#).
- [27] Chi-Fang Chen and Fernando G.S.L. Brandão, *Concentration for Trotter error*, 2021, [arXiv:2111.05324](#).
- [28] Chi-Fang Chen, Hsin-Yuan Huang, Richard Kueng, and Joel A. Tropp, *Concentration for random product formulas*, PRX Quantum **2** (2021), 040305, [arXiv:2008.11751](#).
- [29] Andrew M. Childs, Jiaqi Leng, Tongyang Li, Jin-Peng Liu, and Chenyi Zhang, *Quantum simulation of real-space dynamics*, 2022, [arXiv:2203.17006](#).
- [30] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su, *Toward the first quantum simulation with quantum speedup*, Proceedings of the National Academy of Sciences **115** (2018), no. 38, 9456–9461, [arXiv:1711.10980](#).
- [31] Andrew M. Childs and Yuan Su, *Nearly optimal lattice simulation by product formulas*, Physical Review Letters **123** (2019), 050503, [arXiv:1901.00564](#).
- [32] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu, *Theory of Trotter error with commutator scaling*, Physical Review X **11** (2021), 011020, [arXiv:1912.08854](#).
- [33] Laura Clinton, Toby Cubitt, Brian Flynn, Filippo Maria Gambetta, Joel Klassen, Ashley Montanaro, Stephen Piddock, Raul A. Santos, and Evan Sheridan, *Towards near-term quantum simulation of materials*, 2022, [arXiv:2205.15256](#).
- [34] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein, *Introduction to algorithms*, MIT press, 2022.
- [35] M. Dalmonte, G. Pupillo, and P. Zoller, *One-dimensional quantum liquids with power-law interactions: The luttinger staircase*, Physical Review Letters **105** (2010), 140401, [arXiv:1004.5035](#).
- [36] Xiaolong Deng, Guido Masella, Guido Pupillo, and Luis Santos, *Universal algebraic growth of entanglement entropy in many-body localized systems with power-law interactions*, Physical Review Letters **125** (2020), 010401, [arXiv:1912.08131](#).
- [37] Yulong Dong, Lin Lin, and Yu Tong, *Ground-state preparation and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices*, PRX Quantum **3** (2022), 040305, [arXiv:2204.05955](#).
- [38] J. S. Douglas, H. Habibian, C.-L. Hung, A. V. Gorshkov, H. J. Kimble, and D. E. Chang, *Quantum many-body models with cold atoms coupled to photonic crystals*, Nature Photonics **9** (2015), 326, [arXiv:1312.2435](#).
- [39] Richard P. Feynman, *Simulating physics with computers*, International Journal of Theoretical Physics **21** (1982), no. 6-7, 467–488.
- [40] Yimin Ge, Jordi Tura, and J. Ignacio Cirac, *Faster ground state preparation and high-precision ground energy estimation with fewer qubits*, Journal of Mathematical Physics **60** (2019), no. 2, 022202, [arXiv:1712.03193](#).

- [41] Sevag Gharibian and François Le Gall, *Dequantizing the quantum singular value transformation: Hardness and applications to quantum chemistry and the quantum PCP conjecture*, Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, p. 1932, 2022, [arXiv:2111.09079](#).
- [42] Mariano Giaquinta and Giuseppe Modica, *Mathematical analysis: An introduction to functions of several variables*, Springer Science & Business Media, 2010.
- [43] Mariano Giaquinta and Giuseppe Modica, *Mathematical analysis: Foundations and advanced techniques for functions of several variables*, Springer Science & Business Media, 2011.
- [44] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe, *Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, pp. 193–204, 2019, [arXiv:1806.01838](#).
- [45] A. V. Gorshkov, P. Rabl, G. Pupillo, A. Micheli, P. Zoller, M. D. Lukin, and H. P. Büchler, *Suppression of inelastic collisions between polar molecules with a repulsive shield*, Physical Review Letters **101** (2008), 073201.
- [46] Lars Grasedyck and Wolfgang Hackbusch, *Construction and arithmetics of H-matrices*, Computing **70** (2003), no. 4, 295–334.
- [47] L Greengard and V Rokhlin, *A fast algorithm for particle simulations*, Journal of Computational Physics **73** (1987), no. 2, 325–348.
- [48] Shouzhen Gu, Rolando D. Somma, and Burak Şahinoğlu, *Fast-forwarding quantum evolution*, Quantum **5** (2021), 577, [arXiv:2105.07304](#).
- [49] Jeongwan Haah, Matthew B. Hastings, Robin Kothari, and Guang Hao Low, *Quantum algorithm for simulating real time evolution of lattice Hamiltonians*, Proceedings of the 59th IEEE Symposium on Foundations of Computer Science, pp. 350–360, 2018, [arXiv:1801.03922](#).
- [50] W. Hackbusch, *A sparse matrix arithmetic based on H-matrices. Part I: Introduction to H-matrices*, Computing **62** (1999), no. 2, 89–108.
- [51] Stuart Hadfield, *On the representation of boolean and real functions as Hamiltonians for quantum computing*, ACM Transactions on Quantum Computing **2** (2021), no. 4, 1–21, [arXiv:1804.09130](#).
- [52] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd, *Quantum algorithm for linear systems of equations*, Physical Review Letters **103** (2009), no. 15, 150502, [arXiv:0811.3171](#).
- [53] Takuya Hatomura, *State-dependent error bound for digital quantum simulation of driven systems*, Physical Review A **105** (2022), L050601, [arXiv:2201.04835](#).
- [54] Trygve Helgaker, Poul Jørgensen, and Jeppe Olsen, *Molecular electronic-structure theory*, John Wiley & Sons, 2013.
- [55] Markus Heyl, Philipp Hauke, and Peter Zoller, *Quantum localization bounds Trotter errors in digital quantum simulation*, Science Advances **5** (2019), no. 4, eaau8342, [arXiv:1806.11123](#).
- [56] Roger A. Horn and Charles R. Johnson, *Matrix analysis*, Cambridge university press, 2012.

- [57] Hsin-Yuan Huang, Yu Tong, Di Fang, and Yuan Su, *Learning many-body Hamiltonians with Heisenberg-limited scaling*, 2022, [arXiv:2210.03030](#).
- [58] Aleksei V. Ivanov, Christoph Snderhauf, Nicole Holzmann, Tom Ellaby, Rachel N. Kerber, Glenn Jones, and Joan Camps, *Quantum computation for periodic solids in second quantization*, 2022, [arXiv:2210.02403](#).
- [59] Yifan Jia and Michael M. Wolf, *Hay from the haystack: explicit examples of exponential quantum circuit complexity*, 2022, [arXiv:2205.06977](#).
- [60] N Cody Jones, James D Whitfield, Peter L McMahon, Man-Hong Yung, Rodney Van Meter, Alán Aspuru-Guzik, and Yoshihisa Yamamoto, *Faster quantum chemistry simulation on fault-tolerant quantum computers*, *New Journal of Physics* **14** (2012), no. 11, 115023, [arXiv:1204.0567](#).
- [61] Ivan Kassal, Stephen P. Jordan, Peter J. Love, Masoud Mohseni, and Alán Aspuru-Guzik, *Polynomial-time quantum algorithm for the simulation of chemical dynamics*, *Proceedings of the National Academy of Sciences* **105** (2008), no. 48, 18681–18686, [arXiv:0801.2986](#).
- [62] K. Kim, S. Korenblit, R. Islam, E. E. Edwards, M.-S. Chang, C. Noh, H. Carmichael, G.-D. Lin, L.-M. Duan, C. C. Joseph Wang, J. K. Freericks, and C. Monroe, *Quantum simulation of the transverse Ising model with trapped ions*, *New Journal of Physics* **13** (2011), no. 10, 105003.
- [63] A Yu Kitaev, *Quantum computations: algorithms and error correction*, *Russian Mathematical Surveys* **52** (1997), no. 6, 1191–1249.
- [64] Ian D. Kivlichan, Craig Gidney, Dominic W. Berry, Nathan Wiebe, Jarrod McClean, Wei Sun, Zhang Jiang, Nicholas Rubin, Austin Fowler, Alán Aspuru-Guzik, Hartmut Neven, and Ryan Babbush, *Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via Trotterization*, *Quantum* **4** (2020), 296, [arXiv:1902.10673](#).
- [65] 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*, *Physical Review Letters* **120** (2018), 110501, [arXiv:1711.04789](#).
- [66] Ian D Kivlichan, Nathan Wiebe, Ryan Babbush, and Alán Aspuru-Guzik, *Bounding the costs of quantum simulation of many-body physics in real space*, *Journal of Physics A: Mathematical and Theoretical* **50** (2017), no. 30, 305301, [arXiv:1608.05696](#).
- [67] E. Knill, *Approximation by quantum circuits*, 1995, [arXiv:quant-ph/9508006](#).
- [68] Lucas Kocia, Fernando A. Calderon-Vargas, Matthew D. Grace, Alicia B. Magann, James B. Larsen, Andrew D. Baczewski, and Mohan Sarovar, *Digital adiabatic state preparation error scales better than you might expect*, 2022, [arXiv:2209.06242](#).
- [69] Thomas Koffel, M. Lewenstein, and Luca Tagliacozzo, *Entanglement entropy for the long-range Ising chain in a transverse field*, *Physical Review Letters* **109** (2012), 267203, [arXiv:1207.3957](#).
- [70] Efehan Kökcü, Daan Camps, Lindsay Bassman, J. K. Freericks, Wibe A. de Jong, Roel Van Beeumen, and Alexander F. Kemper, *Algebraic compression of quantum circuits for Hamiltonian evolution*, *Physical Review A* **105** (2022), 032420, [arXiv:2108.03282](#).

- [71] Efehan Kökcü, Thomas Steckmann, Yan Wang, J. K. Freericks, Eugene F. Dumitrescu, and Alexander F. Kemper, *Fixed depth Hamiltonian simulation via Cartan decomposition*, Physical Review Letters **129** (2022), 070501, [arXiv:2104.00728](#).
- [72] Robin Kothari, *Efficient algorithms in quantum query complexity*, Ph.D. thesis, University of Waterloo, 2014.
- [73] Tomotaka Kuwahara, Álvaro M Alhambra, and Anurag Anshu, *Improved thermal area law and quasilinear time algorithm for quantum Gibbs states*, Physical Review X **11** (2021), no. 1, 011047, [arXiv:2007.11174](#).
- [74] David Layden, *First-order Trotter error from a second-order perspective*, Physical Review Letters **128** (2022), 210501, [arXiv:2107.08032](#).
- [75] 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** (2021), 030305, [arXiv:2011.03494](#).
- [76] Lin Lin and Yu Tong, *Near-optimal ground state preparation*, Quantum **4** (2020), 372, [arXiv:2002.12508](#).
- [77] Lin Lin and Yu Tong, *Low-rank representation of tensor network operators with long-range pairwise interactions*, SIAM Journal on Scientific Computing **43** (2021), no. 1, A164–A192, [arXiv:1909.02206](#).
- [78] Lin Lin and Yu Tong, *Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers*, PRX Quantum **3** (2022), 010318, [arXiv:2102.11340](#).
- [79] Seth Lloyd, *Universal quantum simulators*, Science **273** (1996), no. 5278, 1073–1078.
- [80] Ignacio Loaiza, Alireza Marefat Khah, Nathan Wiebe, and Artur F. Izmaylov, *Reducing molecular electronic Hamiltonian simulation cost for linear combination of unitaries approaches*, 2022, [arXiv:2208.08272](#).
- [81] Guang Hao Low, *Hamiltonian simulation with nearly optimal dependence on spectral norm*, Proceedings of the 51th ACM Symposium on Theory of Computing, pp. 491–502, ACM, 2019, [arXiv:1807.03967](#).
- [82] Guang Hao Low and Isaac L. Chuang, *Hamiltonian simulation by uniform spectral amplification*, 2017, [arXiv:1707.05391](#).
- [83] Guang Hao Low and Isaac L. Chuang, *Optimal Hamiltonian simulation by quantum signal processing*, Physical Review Letters **118** (2017), 010501, [arXiv:1606.02685](#).
- [84] Guang Hao Low and Isaac L. Chuang, *Hamiltonian simulation by qubitization*, Quantum **3** (2019), 163, [arXiv:1610.06546](#).
- [85] Guang Hao Low, Vadym Kliuchnikov, and Luke Schaeffer, *Trading T-gates for dirty qubits in state preparation and unitary synthesis*, 2018, [arXiv:1812.00954](#).
- [86] Guang Hao Low, Vadym Kliuchnikov, and Nathan Wiebe, *Well-conditioned multiproduct Hamiltonian simulation*, 2019, [arXiv:1907.11679](#).

- [87] Guang Hao Low and Nathan Wiebe, *Hamiltonian simulation in the interaction picture*, 2018, [arXiv:1805.00675](#).
- [88] Yuri Manin, *Vychislimoe i nevychislimoe (computable and noncomputable)*, Moscow: Sov, 1980.
- [89] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang, *Grand unification of quantum algorithms*, PRX Quantum **2** (2021), 040203, [arXiv:2105.02859](#).
- [90] Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C. Benjamin, and Xiao Yuan, *Quantum computational chemistry*, Reviews of Modern Physics **92** (2020), no. 1, 015003, [arXiv:1808.10402](#).
- [91] Richard Meister, Simon C. Benjamin, and Earl T. Campbell, *Tailoring term truncations for electronic structure calculations using a Linear Combination of Unitaries*, Quantum **6** (2022), 637, [arXiv:2007.11624](#).
- [92] A. Micheli, G. Pupillo, H. P. Büchler, and P. Zoller, *Cold polar molecules in two-dimensional traps: Tailoring interactions with external fields for novel quantum phases*, Physical Review A **76** (2007), 043604, [arXiv:quant-ph/0703031](#).
- [93] Mario Motta and Julia E. Rice, *Emerging quantum computing algorithms for quantum chemistry*, WIREs Computational Molecular Science **12** (2022), no. 3, e1580, [arXiv:2109.02873](#).
- [94] 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 Information **7** (2021), no. 1, 83, [arXiv:1808.02625](#).
- [95] Christopher R. Myers, Cyrus J. Umrigar, James P. Sethna, and John D. Morgan, *Fock's expansion, Kato's cusp conditions, and the exponential ansatz*, Physical Review A **44** (1991), 5537–5546.
- [96] Quynh T. Nguyen, Bobak T. Kiani, and Seth Lloyd, *Quantum algorithm for dense and full-rank kernels using hierarchical matrices*, 2022, [arXiv:2201.11329](#).
- [97] Bryan O’Gorman, William J. Huggins, Eleanor G. Rieffel, and K. Birgitta Whaley, *Generalized swap networks for near-term quantum computing*, 2019, [arXiv:1905.05118](#).
- [98] 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*, Quantum Information and Computation **15** (2015), no. 5-6, 361–384, [arXiv:1406.4920](#).
- [99] David Poulin, Alexei Y Kitaev, Damian Steiger, Matthew Hastings, and Matthias Troyer, *Quantum algorithm for spectral measurement with lower gate count*, Physical Review Letters **121** (2018), no. 1, 010501, [arXiv:1711.11025](#).
- [100] David Poulin and Pawel Wocjan, *Preparing ground states of quantum many-body systems on a quantum computer*, Physical Review Letters **102** (2009), 130503, [arXiv:0809.2705](#).
- [101] Abhishek Rajput, Alessandro Roggero, and Nathan Wiebe, *Hybridized methods for quantum simulation in the interaction picture*, Quantum **6** (2022), 780, [arXiv:2109.03308](#).

- [102] Patrick Rall, *Faster coherent quantum algorithms for phase, energy, and amplitude estimation*, Quantum **5** (2021), 566, [arXiv:2103.09717](#).
- [103] Tim Roughgarden, *Algorithms illuminated*, Soundlikeyourself publishing, 2017.
- [104] Nicholas C. Rubin, Joonho Lee, and Ryan Babbush, *Compressing many-body fermion operators under unitary constraints*, 2021, [arXiv:2109.05010](#).
- [105] M. Saffman, T. G. Walker, and K. Mølmer, *Quantum information with Rydberg atoms*, Reviews of Modern Physics **82** (2010), 2313–2363, [arXiv:0909.4777](#).
- [106] Burak Şahinoğlu and Rolando D. Somma, *Hamiltonian simulation in the low-energy subspace*, npj Quantum Information **7** (2021), no. 1, 119, [arXiv:2006.02660](#).
- [107] Alexander F. Shaw, Pavel Lougovski, Jesse R. Stryker, and Nathan Wiebe, *Quantum algorithms for simulating the lattice Schwinger model*, Quantum **4** (2020), 306, [arXiv:2002.11146](#).
- [108] V.V. Shende, S.S. Bullock, and I.L. Markov, *Synthesis of quantum-logic circuits*, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **25** (2006), no. 6, 1000–1010, [arXiv:quant-ph/0406176](#).
- [109] Mark Steudtner and Stephanie Wehner, *Estimating exact energies in quantum simulation without Toffoli gates*, Physical Review A **101** (2020), 052329, [arXiv:1911.00416](#).
- [110] Yuan Su, Dominic W. Berry, Nathan Wiebe, Nicholas Rubin, and Ryan Babbush, *Fault-tolerant quantum simulations of chemistry in first quantization*, PRX Quantum **2** (2021), 040332, [arXiv:2105.12767](#).
- [111] Yuan Su, Hsin-Yuan Huang, and Earl T. Campbell, *Nearly tight Trotterization of interacting electrons*, Quantum **5** (2021), 495, [arXiv:2012.09194](#).
- [112] Yu Tong, Victor V. Albert, Jarrod R. McClean, John Preskill, and Yuan Su, *Provably accurate simulation of gauge theories and bosonic systems*, Quantum **6** (2022), 816, [arXiv:2110.06942](#).
- [113] Minh C. Tran, Andrew Y. Guo, Yuan Su, James R. Garrison, Zachary Eldredge, Michael Foss-Feig, Andrew M. Childs, and Alexey V. Gorshkov, *Locality and digital quantum simulation of power-law interactions*, Physical Review X **9** (2019), 031006, [arXiv:1808.05225](#).
- [114] F Verstraete and J I Cirac, *Mapping local Hamiltonians of fermions to local Hamiltonians of spins*, Journal of Statistical Mechanics: Theory and Experiment **2005** (2005), no. 09, P09012, [arXiv:cond-mat/0508353](#).
- [115] Vera von Burg, Guang Hao Low, Thomas Häner, Damian S. Steiger, Markus Reiher, Martin Roetteler, and Matthias Troyer, *Quantum computing enhanced computational catalysis*, Physical Review Research **3** (2021), 033055, [arXiv:2007.14460](#).
- [116] Kianna Wan, Mario Berta, and Earl T. Campbell, *Randomized quantum algorithm for statistical phase estimation*, Physical Review Letters **129** (2022), 030503, [arXiv:2110.12071](#).
- [117] Jonathan Welch, Daniel Greenbaum, Sarah Mostame, and Alan Aspuru-Guzik, *Efficient quantum circuits for diagonal unitaries without ancillas*, New Journal of Physics **16** (2014), no. 3, 033040, [arXiv:1306.3991](#).

- [118] Shenglong Xu, Leonard Susskind, Yuan Su, and Brian Swingle, *A sparse model of quantum holography*, 2020, [arXiv:2008.02303](#).
- [119] Bo Yan, Steven A. Moses, Bryce Gadway, Jacob P. Covey, Kaden R. A. Hazzard, Ana Maria Rey, Deborah S. Jin, and Jun Ye, *Observation of dipolar spin-exchange interactions with lattice-confined polar molecules*, *Nature* **501** (2013), 521.
- [120] Qi Zhao and Xiao Yuan, *Exploiting anticommutation in Hamiltonian simulation*, *Quantum* **5** (2021), 534, [arXiv:2103.07988](#).
- [121] Qi Zhao, You Zhou, Alexander F. Shaw, Tongyang Li, and Andrew M. Childs, *Hamiltonian simulation with random inputs*, 2021, [arXiv:2111.04773](#).