# Toward the first quantum simulation with quantum speedup 

Andrew M. Childs, ${ }^{1,2,3, *}$ Dmitri Maslov, ${ }^{2,3,4}$ Yunseong Nam ${ }^{2,3,5}$, Neil J. Ross ${ }^{2,3,6}$, and Yuan $\mathrm{Su}^{1,2,3}$<br>${ }^{1}$ Department of Computer Science, University of Maryland<br>${ }^{2}$ Institute for Advanced Computer Studies, University of Maryland<br>${ }^{3}$ Joint Center for Quantum Information and Computer Science, University of Maryland<br>${ }^{4}$ National Science Foundation<br>${ }^{5}$ IonQ, Inc.<br>${ }^{6}$ Department of Mathematics and Statistics, Dalhousie University


#### Abstract

With quantum computers of significant size now on the horizon, we should understand how to best exploit their initially limited abilities. To this end, we aim to identify a practical problem that is beyond the reach of current classical computers, but that requires the fewest resources for a quantum computer. We consider quantum simulation of spin systems, which could be applied to understand condensed matter phenomena. We synthesize explicit circuits for three leading quantum simulation algorithms, employing diverse techniques to tighten error bounds and optimize circuit implementations. Quantum signal processing appears to be preferred among algorithms with rigorous performance guarantees, whereas higher-order product formulas prevail if empirical error estimates suffice. Our circuits are orders of magnitude smaller than those for the simplest classically-infeasible instances of factoring and quantum chemistry.


## 1 Introduction

While a scalable quantum computer remains a long-term goal, recent experimental progress suggests that devices capable of outperforming classical computers will soon be available [10, 21, 28, 38, 73, 78]. Multiple groups have already developed programmable devices with several qubits and twoqubit gate fidelities around $98 \%$ [49], and similar devices with around 50 qubits are under active development. While the error rates of these early machines severely limit the total number of gates that can be reliably performed, future improvements should lead to machines with more qubits and more reliable gates. This raises the exciting possibility of solving practical problems that are beyond the reach of classical computation. Such an outcome would be a landmark in the development of quantum computers and would begin an era in which they serve not only as testbeds for science, but as practical computing machines.

Reaching this goal will require not only significant experimental advances, but also careful quantum algorithm design and implementation. Here we address the latter issue by developing explicit circuits, and thereby producing concrete resource estimates, for practical quantum computations that can outperform classical computers. Through this work, we aim to identify applications for small quantum computers that help to motivate the significant investment required to develop scalable, fault-tolerant quantum computers.

There has been considerable previous research on compiling quantum algorithms into explicit circuits (see Appendix A for more detail). However, to the best of our knowledge, none of these
studies aimed to identify minimal examples of super-classical quantum computation, and typical resource counts were high. Our work is also distinct from recent work on quantum computational supremacy [35], where the goal is merely to accomplish a super-classical task, regardless of its practicality. Instead, we aim to pave the way toward practical quantum computations (which may not be far beyond the threshold for supremacy).

Arguably, the most natural application of quantum computers is to the problem of simulating quantum dynamics [30]. Quantum computers can simulate a wide variety of quantum systems, including fermionic lattice models [77], quantum chemistry [76], and quantum field theories [40]. However, simulations of spin systems with local interactions likely have less overhead, so we focus on them as an early candidate for practical quantum simulation. While analog simulation may be easier to realize in the short term, we focus on digital simulation for its greater flexibility and the prospect of invoking fault tolerance.

Efficient quantum algorithms for simulating quantum dynamics have been known for over two decades [50]. Recent work has led to algorithms with significantly improved asymptotic performance as a function of various parameters such as the evolution time and the allowed simulation error $[11,12,14,51,52]$. Our work investigates whether these alternative algorithms can be advantageous for simulations of relatively small systems, and aims to lay the groundwork for the first practical application of quantum computers.

## 2 Target system

To produce concrete benchmarks, we focus on a specific simulation task. Specifically, we consider a one-dimensional nearest-neighbor Heisenberg model with a random magnetic field in the $z$ direction. This model is described by the Hamiltonian

$$
\begin{equation*}
\sum_{j=1}^{n}\left(\vec{\sigma}_{j} \cdot \vec{\sigma}_{j+1}+h_{j} \sigma_{j}^{z}\right) \tag{1}
\end{equation*}
$$

where $\vec{\sigma}_{j}=\left(\sigma_{j}^{x}, \sigma_{j}^{y}, \sigma_{j}^{z}\right)$ denotes a vector of Pauli $x, y$, and $z$ matrices on qubit $j$. We impose periodic boundary conditions (i.e., $\vec{\sigma}_{n+1}=\vec{\sigma}_{1}$ ), and $h_{j} \in[-h, h]$ is chosen uniformly at random. The parameter $h$ characterizes the strength of the disorder.

This Hamiltonian has been considered in recent studies of self-thermalization and many-body localization (see Appendix B for more detail). Despite intensive investigation, the details of a transition between thermal and localized phases remain poorly understood. A major challenge is the difficulty of simulating quantum systems with classical computers; indeed, the most extensive numerical study we are aware of was restricted to at most 22 spins [54].

Hamiltonian simulation can efficiently access any feature that could be observed experimentally (and more), and there are several proposals for exploring self-thermalization by simulating dynamics [67, 68, 70]. Since all of these approaches involve only very simple state preparations and measurements, we focus on the cost of simulating dynamics. We consider evolution times comparable to the number of spins, since the system must evolve for this long for self-thermalization to take place (or even for information to propagate across the system, owing to the Lieb-Robinson bound).

Specifically, we produce gate counts for simulations with $h=1$, evolution time $t=n$ (the number of spins in the system), and overall accuracy $\epsilon=10^{-3}$. These explicit choices help us to focus on the system-size dependence of quantum simulation algorithms. This is a key consideration for practical applications, yet it has been deemphasized in the literature on sparse Hamiltonian simulation.

| Algorithm | Ref. | Gate complexity $(t, \epsilon)$ | Gate complexity $(n)$ |
| :--- | :---: | :--- | :--- |
| Product formula (PF), 1st order | $[50]$ | $O\left(t^{2} / \epsilon\right)$ | $O\left(n^{5}\right)$ |
| Product formula (PF), $(2 k)$ th order | $[11]$ | $O\left(5^{2 k} t^{1+1 / 2 k} / \epsilon^{1 / 2 k}\right)$ | $O\left(5^{2 k} n^{3+1 / k}\right)$ |
| Quantum walk | $[12]$ | $O(t / \sqrt{\epsilon})$ | $O\left(n^{4} \log n\right)$ |
| Fractional-query simulation | $[13]$ | $O\left(t \frac{\log ^{2}(t / \epsilon)}{\log \log ^{2}(t / \epsilon)}\right)$ | $O\left(n^{4} \frac{\log n}{\log \log n}\right)$ |
| Taylor series (TS) | $[14]$ | $O\left(t \frac{\log ^{2}(t / \epsilon)}{\log \log (t / \epsilon)}\right)$ | $O\left(n^{3} \frac{\log ^{2} n}{\log \log n}\right)$ |
| Linear combination of quantum walk steps | $[15]$ | $O\left(t \frac{\log ^{3.5}(t / \epsilon)}{\log \log (t / \epsilon)}\right)$ | $O\left(n^{4} \frac{\log n}{\log \log n}\right)$ |
| Quantum signal processing (QSP) | $[51]$ | $O(t+\log (1 / \epsilon))$ | $O\left(n^{3} \log n\right)$ |

Table 1: Previously-established asymptotic gate complexities of quantum simulation algorithms as a function of the simulation time $t$, allowed error $\epsilon$, and the system size $n$ for a one-dimensional nearest-neighbor spin model as in (1) with $t=n$ and fixed $\epsilon$.

## 3 Implementations

There are many distinct quantum algorithms for Hamiltonian simulation, some of which are summarized in Table 1. We implement algorithms based on high-order product formulas (PF, introduced in Section C.1) [11], direct application of the Taylor series (TS, Section C.2) [14], and the recent quantum signal processing method (QSP, Section C.3) [51]. We expect these to be among the most efficient approaches to digital quantum simulation. In particular, approaches based on quantum walk $[12,15]$ appear to incur greater overhead (as discussed in Appendix D).

To produce concrete circuits, we implement quantum simulation algorithms in a quantum circuit description language called Quipper [32] (see Appendix E for more details). Wherever possible, we tighten the analysis of algorithm parameters and manually optimize the implementation. We also process all circuits using an automated tool we developed for large-scale quantum circuit optimization [57]. Our implementation is available in a public repository [26].

We express our circuits over the set of two-qubit CNOT gates, single-qubit Clifford gates, and single-qubit $z$ rotations $R_{z}(\theta):=\exp \left(-i \sigma^{z} \theta / 2\right)$ for $\theta \in \mathbb{R}$. Such gates can be directly implemented at the physical level with both trapped ions [28] and superconducting circuits [21, 38]. In both technologies, two-qubit gates take longer to perform and incur more error than single-qubit gates. Thus, the CNOT count is a useful figure of merit for assessing the cost of physical-level circuits on a universal device. We also produce Clifford $+T$ circuits using optimal circuit synthesis [66] so that we can count $T$ gates, which are typically the most expensive gates for fault-tolerant computation.

Our analysis ignores many practical details, such as architectural constraints, instead aiming to give a broad overview of potential implementation costs that can be refined for specific systems. When counting qubits, we assume that measured ancillas can be reused later.

### 3.1 Product formula algorithm

The product formula (PF) approach approximates the exponential of a sum of operators by a product of exponentials of the individual operators. The asymptotic complexity of this approach can be improved with higher-order Suzuki formulas [74]. By splitting the evolution into $r$ segments and making $r$ sufficiently large, we can ensure that the simulation is arbitrarily precise. The main challenge in making these algorithms concrete is to choose an explicit $r$ that ensures some desired upper bound on the error. Appendix F gives a detailed description of these implementation details.

We present two bounds, which we call the analytic and minimized bounds, that slightly strengthen previous analysis [11]. However, bounds of this type are far from tight [5, 63, 64]. Thus, we develop
an improved bound that exploits commutation relations among terms in the target Hamiltonian. For a one-dimensional system of $n$ spins with nearest-neighbor couplings, evolving for time $t=n$, this commutator bound improves the asymptotic complexity of the $(2 k)$ th-order PF algorithm from $O\left(n^{3+1 / k}\right)$ to $O\left(n^{3+2 /(2 k+1)}\right)$ while also significantly improving the leading constant.

Naive computation of the commutator bound takes time $O\left(n^{2 k+1}\right)$, which can be prohibitive even for small $n$. To make this approach practical, we develop techniques that exploit the structure of the Hamiltonian to compute the commutator bound in closed form. We explicitly evaluate this bound for the first-, second-, and fourth-order product formulas.

Unfortunately, even the commutator bound can be very loose. To address this, we report empirical error estimates by extrapolating the error seen in direct classical simulations of small instances (as also explored in previous work on simulating many-body dynamics [63] and quantum chemistry [5, 64]). While these empirical bounds do not provide rigorous guarantees on the simulation error, they may nevertheless be useful in practice, and they improve the cost of PF algorithms by several orders of magnitude.

### 3.2 Taylor series algorithm

The Taylor series (TS) algorithm directly implements the (truncated) Taylor series of the evolution operator for a carefully-chosen constant time using a procedure for implementing linear combinations of unitary operations [14]. This segment is then simply repeated until the entire evolution time has been simulated. The circuit for a segment is built using three subroutines: a state preparation procedure, a reflection about the $|0\rangle$ state, and an operation denoted select $(V)$ (discussed further below). Our implementation of the TS algorithm (described in detail in Appendix G) also includes a concrete error analysis that establishes rigorous, non-asymptotic bounds on the simulation parameters.

The aforementioned select $(V)$ operation applies a unitary $V_{j}$ conditioned on a control register being in the state $|j\rangle$, for $j \in\{1, \ldots, \Gamma\}$. We develop an improved implementation of this operation by designing an optimized walk on a binary tree, saving a factor of about $\log _{2} \Gamma$ in the gate count. For our simulations of systems with 10 to 100 spins, this reduces CNOT and $T$ gate counts over a naive implementation by a factor of between 5 and 9 , significantly improving the overall complexity. Furthermore, the cost of our select $(V)$ implementation meets a previously-established asymptotic lower bound [56]. This improvement may be more broadly applied to any algorithm using the select $(V)$ procedure, such as others based on linear combinations of unitaries.

### 3.3 Quantum signal processing algorithm

The quantum signal processing (QSP) algorithm of Low and Chuang [51, 52] effectively implements a linear combination of unitary operators by a different mechanism. This algorithm applies a sequence of operations called phased iterates that manifest each eigenvalue of the Hamiltonian as a rotation acting on an ancilla qubit. By carefully choosing a sequence of rotation angles for that qubit, we induce the desired evolution.

The circuit for each phased iterate is built from similar subroutines as the TS algorithm. However, computing the $M$ rotation angles for the phased iterates requires finding the roots of a polynomial of degree $2 M$, and these roots must be determined to high precision. Because of these challenges, we were unable to compute the parameters of the algorithm explicitly except in very small instances. Instead, we produced estimates of the gate count (but not a complete implementation) by synthesizing a version of the algorithm with placeholder values of the parameters.


Figure 1: Gate counts for optimized implementations of the PF algorithm (using the fourth-order formula with commutator bound and the better of the fourth- or sixth-order formula with empirical error bound), the TS algorithm, and the QSP algorithm (using the segmented version with analytic error bound and the non-segmented version with empirical Jacobi-Anger error bound) for system sizes between 10 and 100. Left: CNOT gates for Clifford $+R_{z}$ circuits. Right: $T$ gates for Clifford $+T$ circuits.

One way to alleviate this problem is to consider a segmented implementation of the QSP algorithm. In this approach, we divide the evolution time into $r$ segments, each of which is sufficiently short that the classical preprocessing is tractable. Since the optimality of the QSP approach to Hamiltonian simulation relies essentially on simulating the entire evolution as a single segment, the segmented approach has higher asymptotic complexity. However, it allows us to develop a complete implementation, and the overhead for moderate values of $n$ is not too high.

For the full version of the algorithm, we consider an empirical error bound on the Jacobi-Anger expansion, giving a modest improvement. Numerical evidence suggests that the additional savings from an empirical error bound for the overall algorithm would not be significant. For the segmented version of the algorithm, we instead used an analytic error bound so that the algorithm remains rigorous (and because an empirical Jacobi-Anger error bound did not give much improvement in that case).

Appendix H discusses our implementation of QSP algorithms in detail.

## 4 Results

Figure 1 compares gate counts for the PF algorithm (with commutator and empirical error bounds), the TS algorithm, and the QSP algorithm (in both segmented and non-segmented versions). The TS algorithm uses significantly more qubits than the QSP algorithm (as shown in Figure 2) while also requiring more gates, so the latter is clearly preferred. In contrast, the QSP algorithm has only slightly greater space requirements than the PF algorithm.

Surprisingly, despite being more involved, the QSP algorithm outperforms the rigorously-bounded PF algorithm even for small system sizes. In particular, among the rigorously-analyzed algorithms, the segmented QSP algorithm has the best performance, improving over the PF algorithm by about an order of magnitude for CNOT count and by almost two orders of magnitude for $T$ count.

Empirical error bounds improve the performance of the PF algorithm by two to three orders of magnitude, making it the preferred approach if rigorous performance guarantees are not required. For the CNOT count, the empirical PF algorithm improves over the full QSP algorithm by about an order of magnitude. The advantage in the $T$ count is less significant, but still indicates that the


Figure 2: Number of qubits used by the PF, TS, and QSP algorithms.


Figure 3: Total gate counts in the Clifford $+R_{z}$ basis for product formula algorithms using the minimized (left), commutator (center), and empirical (right) bounds, for system sizes between 13 and 500 .

PF algorithm is dominant, especially considering its lower qubit count.
Although we expected that higher-order product formulas would not be advantageous for small system sizes, we find that the fourth- and sixth-order formulas had the best performance for our benchmark system with tens to hundreds of qubits, as shown in Figure 3. The fourth-order formula with commutator bound gives the best available PF algorithm with a rigorous performance guarantee. Using empirical error bounds, the sixth-order formula outperforms the fourth-order formula for systems of about 30 or more qubits, making the former the method of choice for simulations just beyond the reach of classical computers (again, provided a heuristic error bound can be tolerated). These results suggest that higher-order formulas may be advantageous for other quantum simulations, such as those for quantum chemistry, even though they have not usually been considered [62, 64].

For a system of 50 qubits - which is presumably close to the limits of direct classical simulation for circuits such as ours [34] ${ }^{1}$ - the TS algorithm uses 171 qubits and the QSP algorithm uses 67, whereas the PF algorithm uses only 50. At this size, the segmented QSP algorithm is the best rigorously-analyzed approach, using about $1.8 \times 10^{8}$ CNOT gates (over the set of Clifford $+R_{z}$ gates) and $2.4 \times 10^{9} T$ gates (over the set of Clifford $+T$ gates). Using the empirical error bound, the PF algorithm uses about $3 \times 10^{6}$ CNOTs and $1.8 \times 10^{8} T \mathrm{~s}$ (over Clifford $+R_{z}$ and Clifford $+T$,

[^0]respectively).
For comparison, previous estimates of gate counts for factoring, discrete logarithms, and quantum chemistry simulations are significantly larger. First consider factoring a 1024-bit number, which is beyond the factorization of RSA-768 that was achieved classically in 2009 [44]. The best implementation we are aware of uses 3132 qubits and about $5.7 \times 10^{9} T$ gates (when realized over the set of Clifford $+T$ gates) [47]. ${ }^{2}$ Quantum algorithms for classically-hard instances of the elliptic curve discrete logarithm problem have roughly comparable cost [65]. For quantum chemistry, a natural target for a problem just beyond the reach of classical computing is a simulation of FeMoco, the primary cofactor of nitrogenase, an enzyme that catalyzes the nitrogen fixation process. Even for a fairly low-precision simulation, and using non-rigorous estimates of the product formula error, the best implementation we are aware of uses 111 qubits and $1.0 \times 10^{14} T$ gates [64]. Thus it appears that simulation of spin systems is indeed a significantly easier task for near-term practical quantum computation.

For a more detailed discussion of the results, see Appendix I.

## 5 Discussion

The work described in this paper represents progress toward the first genuine application of quantum computers, solving a practical problem that is beyond the reach of classical computation. Of course, our results only represent upper bounds. While we attempted to optimize the implementation wherever possible, it is likely that further improvements can be found, and it is conceivable that another algorithm (or computational task) may offer better performance. Our work establishes a concrete set of benchmarks that we hope can be improved through future studies.

Demonstrations of digital quantum simulation performed to date [ $6,19,48$ ] have been limited in scope, primarily using the first-order formula (except for some limited applications of the secondorder formula $[19,48])$. Our results show that higher-order formulas are useful even for simulations of small systems. In the near term, it could be fruitful to demonstrate the utility of these formulas experimentally. Even relatively small experiments might be able to probe the validity of our empirical error bounds.

We have also identified some avenues for future improvement of quantum simulation algorithms. We saw that rigorous error bounds for product formulas are very loose, even with our newlydeveloped commutator bound. This motivates attempting to prove stronger rigorous error bounds for product formulas. Also, the difficulty of computing the angles needed to perform the QSP algorithm prevents us from taking full advantage of the algorithm in practice, so it would be useful to develop a more efficient classical procedure for specifying these angles.

Further reduction of the gate count could be especially significant if it led to a simulation with sufficiently few gates to be performed without invoking fault tolerance. With our current estimate of millions of CNOT gates for a superclassical simulation, this is likely out of reach at present. However, further improvement could obviate the need for error correction in a system with highly accurate gates, making an early demonstration of superclassical simulation more accessible.

Finally, our work has considered an idealized system, and we hope future work will take more realistic constraints into account. Practical devices will come with architectural constraints, may employ different basic operations than those considered here, may allow parallelization of gates, and will likely require fault tolerance. By incorporating such features, we hope the work begun here will lead to a blueprint for the first practical quantum computation.

[^1]
## Acknowledgments

We thank Zhexuan Gong, Alexey Gorshkov, Guang Hao Low, Chris Monroe, and Nathan Wiebe for helpful discussions.

This work was supported in part by the Army Research Office (MURI award W911NF-16-10349), the Canadian Institute for Advanced Research, and the National Science Foundation (grant 1526380).

This material was partially based on work supported by the National Science Foundation during DM's assignment at the Foundation. Any opinion, finding, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## Appendices

A Related work ..... 10
B Self-thermalization in spin models ..... 10
C Simulation algorithms ..... 11
C. 1 Product formula algorithm ..... 12
C. 2 Taylor series algorithm ..... 13
C. 3 Quantum signal processing algorithm ..... 16
D System-size dependence for other simulation algorithms ..... 18
E Circuit synthesis and optimization ..... 20
E. 1 Overview of algorithm implementations ..... 20
E.1.1 Product formulas ..... 20
E.1.2 Taylor series ..... 21
E.1.3 Quantum signal processing ..... 21
E. 2 Gate counts and synthesis ..... 21
E. 3 Optimization ..... 22
E. 4 Correctness ..... 23
F Product formula implementation details ..... 23
F. 1 Analytic and minimized bounds ..... 24
F. 2 Commutator bounds ..... 26
F.2.1 Abstract commutator bounds ..... 27
F.2.2 Concrete commutator bounds ..... 31
F. 3 Empirical bounds ..... 38
G Taylor series implementation details ..... 38
G. 1 Error analysis ..... 38
G. 2 Failure probability ..... 43
G. 3 Encoding of the control register ..... 43
G. 4 Implementation of $\operatorname{select}(V)$ ..... 43
H Quantum signal processing implementation details ..... 48
H. 1 Failure probability ..... 48
H. 2 Circuit optimizations ..... 49
H. 3 Phase computation and segmented algorithm ..... 49
H. 4 Empirical error bounds ..... 51
I Detailed results ..... 52
I. 1 Product formulas ..... 53
I. 2 Other algorithms and comparisons ..... 55
I. 3 Circuit optimization ..... 56

## A Related work

In this appendix, we discuss related work on quantum algorithms for simulating physics and resource estimates for practical quantum computation.

The idea of simulating physics with quantum computers was suggested by Feynman [30] and others in the 1980s. Since that time, there has been significant progress on the development of algorithms for the basic problem of simulating Hamiltonian dynamics [1, 11-15, 23-25, 50-52]. In addition, many authors have developed methods for using quantum computers to simulate the behavior of specific types of systems. Quantum computers presumably offer an advantage for simulations of any system in which quantum mechanics plays a significant role, including fermionic lattice models [63, 71, 72, 77], quantum chemistry [3-5, 36, 42, 62, 64, 76], and quantum field theories [39-41]. Much of this work has focused on asymptotic analysis, although some work on concrete resource estimates for quantum chemistry simulation is mentioned below.

There has been considerable previous research on compiling quantum algorithms into explicit circuits, including work on the IARPA Quantum Computer Science (QCS) program [37]. The QCS program focused on synthesizing a selection of quantum algorithms and developing optimized implementations over certain models of physical machines that were chosen to describe realistic devices. To the best of our knowledge, none of these studies aimed to construct a minimal example of a super-classical quantum computation, and typical resource counts were high.

In addition, many researchers have developed optimized implementations of Shor's integer factoring algorithm $[8,9,31,33,46,47,60,75]$ and estimated resource requirements for simulating quantum chemistry $[36,64,76]$. As discussed in Section 4, our results suggest that quantum simulation of spin models will be accessible with dramatically fewer computational resources, making this a promising candidate for an early demonstration of practical quantum computation.

The best implementation of the factoring algorithm that we are aware of appears in [47]. That paper does not give explicit resource counts, so we estimate them as follows. The implementation of Shor's algorithm described there uses $4 n^{3}+O\left(n^{2} \log (n)\right)$ gates and $3 n+6 \log (n)+O(1)$ qubits [47, page 10]. However, this count assumes we can perform arbitrary 2 -qubit gates. The dominant contribution to the gate count comes from modular exponentiation (in particular, the cost of the QFT is small), which relies mainly on so-called psuedo-Toffoli gates. Each of these gates is realized with two controlled $-H$ gates and one controlled $-Z$ gate; the former needs two $T$ gates and the latter needs none, so the total number of $T$ gates is approximately $\frac{16}{3} n^{3}$.

Note that it is possible to factor an $n$-bit number using only about $2 n$ qubits [8, 33], but at the expense of a significantly higher gate count. Similarly, gate counts for quantum algorithms the elliptic curve discrete logarithm problem can also be reduced at the cost of using significantly more qubits [22].

## B Self-thermalization in spin models

In this appendix we motivate our choice of a candidate spin system to simulate with an early universal quantum computer, elaborating on the brief discussion in Section 2.

Recently, there has been considerable interest among the condensed matter community in understanding the equilibration of closed quantum systems [58]. While equilibration is normally viewed as a consequence of coupling a system to a bath, a large but isolated quantum system may effectively display features of thermalization through its own unitary dynamics-or it may fail to self-thermalize through a phenomenon known as many-body localization. Despite intense study, the details of phase transitions between localized and thermal phases in various model systems remain
poorly understood. A major challenge is the difficulty of simulating quantum systems with classical computers, which has restricted numerical investigations to systems with fewer than 25 qubits.

To produce concrete benchmarks, we focus on a specific simulation task with the potential for practical applications. As discussed in Section 2, we consider a one-dimensional nearest-neighbor Heisenberg model with a random magnetic field in the $z$ direction, described by the Hamiltonian (1). This Hamiltonian has been studied as a model of self-thermalization and many-body localization $[54,58,59]$, where the most extensive numerical study we are aware of was restricted to at most 22 spins [54]. The nature of the phase transition as a function of $h$ remains unclear [58, Section 6.2] and could be illuminated by larger-scale simulations.

Classical numerics typically involve performing exact diagonalization to evaluate properties of the full spectrum of the Hamiltonian (which is inefficient, and even more costly than simulating the dynamics classically). To apply efficient quantum simulation, we must choose an initial state and final measurement so that the outcomes are informative (say, so they provide useful information about the phase diagram). While more limited than a calculation of the full spectrum, a quantum simulation performed on a universal quantum computer should be able to efficiently extract any information that can be probed by an experiment (in addition to other quantities that could be hard to extract experimentally). Experimental probes of thermalization and many-body localization typically involve preparing a product initial state and performing a final standard-basis measurement to see how the outcomes evolve in time - whether they retain memory of the initial state or approach a thermal distribution. Specific observables include the Hamming distance from an initial classical configuration [70] and the imbalance between occupation of even and odd sites [67]. Another proposal considers performing a spin-echo sequence (which involves simulating evolution backward in time, something that is easy to accomplish using digital simulation) [68].

One might also perform similar tests on a (randomly sampled) eigenstate of the Hamiltonian, using phase estimation [7]. However, this introduces additional overhead. Since we expect that it will be informative (if not decisive) to study a global quench, we focus here on the cost of simulating dynamics.

While we have considered the Heisenberg model (1) for concreteness, there is also interest in exploring many-body localization and thermalization (among other phenomena) for diverse spin systems including Ising $[43,70]$ and XXZ [68] models. Our basic approach to estimating simulation resource requirements would apply to these models essentially unchanged and we expect that similar conclusions about the relative performance of different quantum simulation algorithms would hold.

## C Simulation algorithms

As discussed in Section 3, there are many different quantum algorithms for Hamiltonian simulation, including algorithms based on product formulas [11, 23, 50], discrete-time quantum walks [12, 15, 24,52 ], and linear combinations of unitaries [13-15, 27,51,52]. The asymptotic performance of these algorithms is summarized in Table 1, both as a function of the simulation time $t$ and allowed error $\epsilon$ (as commonly emphasized in the literature on sparse Hamiltonian simulation) and as a function of the number of qubits $n$ (for the target system described in Section 2).

In this paper, we focus on the algorithms based on product formulas (PF) [11], on direct implementation of the Taylor series (TS) [14], and on Low and Chuang's recent quantum signal processing (QSP) approach [51]. We expect these three algorithms to be among the best for simulations of spin systems. Considering the dependence on system size, the algorithm based on quantum walk [12] has worse asymptotic performance than the three methods we consider, and requires a costly computation of trigonometric functions performed in quantum superposition.

The algorithm based on the fractional-query model [13] is conceptually similar to the one that implements the Taylor series [14], but the latter has a streamlined form, resulting in improved asymptotic complexity as a function of $n$. For sparse Hamiltonians, the algorithm based on a linear combination of quantum walk steps [15] has improved query complexity as a function of sparsity over [14], but this improvement is not relevant to local Hamiltonians, and again the gate complexity is higher.

As discussed in Section 2, we focus on a specific type of Hamiltonian that can be described as a sum of $L$ terms, of the form

$$
\begin{equation*}
H=\sum_{\ell=1}^{L} \alpha_{\ell} H_{\ell} \tag{2}
\end{equation*}
$$

where each $H_{\ell}$ is a tensor product of Pauli operators acting nontrivially on at most two out of $n$ qubits. We also assume that all the coefficients are positive real numbers $\alpha_{\ell}>0$, since if $\alpha_{\ell}<0$, we can absorb the negative sign into the definition of $H_{\ell}$. In each of the algorithms described below, our goal is to implement an approximation of the unitary operation $\exp (-i H t)$ for a given time $t \in \mathbb{R}$, up to an allowed error at most $\epsilon>0$. Although for some applications we might want to simulate evolutions for negative times (e.g., to implement the spin-echo sequence proposed in [68]), we can absorb this into the sign of the Hamiltonian, so we can consider $t>0$ without loss of generality.

In the remainder of this appendix, we review the three algorithms we consider, emphasizing aspects relevant to our implementations. We consider the PF algorithm in Section C.1, the TS algorithm in Section C.2, and the QSP algorithm in Section C.3. We briefly discuss the system size dependence of other quantum simulation algorithms in Appendix D.

## C. 1 Product formula algorithm

The product formula (PF) approach approximates the exponential of a sum of operators by a product of exponentials of the individual operators. The first-order formula

$$
\begin{equation*}
\left\|\exp \left(-i t \sum_{j=1}^{L} \alpha_{j} H_{j}\right)-\left[\prod_{j=1}^{L} \exp \left(-\frac{i t}{r} \alpha_{j} H_{j}\right)\right]^{r}\right\|=O\left(\frac{(L \Lambda t)^{2}}{r}\right), \tag{3}
\end{equation*}
$$

where $\Lambda:=\max _{j} \alpha_{j}$, underlies the first explicit quantum simulation algorithm [50]. The complexity of quantum simulation can be improved [11, 23] using the ( $2 k$ ) th-order Suzuki formula $S_{2 k}$, defined recursively by [74]

$$
\begin{align*}
S_{2}(\lambda) & :=\prod_{j=1}^{L} \exp \left(\alpha_{j} H_{j} \lambda / 2\right) \prod_{j=L}^{1} \exp \left(\alpha_{j} H_{j} \lambda / 2\right)  \tag{4}\\
S_{2 k}(\lambda) & :=\left[S_{2 k-2}\left(p_{k} \lambda\right)\right]^{2} S_{2 k-2}\left(\left(1-4 p_{k}\right) \lambda\right)\left[S_{2 k-2}\left(p_{k} \lambda\right)\right]^{2}, \tag{5}
\end{align*}
$$

with $p_{k}:=1 /\left(4-4^{1 /(2 k-1)}\right)$ for $k>1$. Using this improved formula, we have [11]

$$
\begin{equation*}
\left\|\exp \left(-i t \sum_{j=1}^{L} \alpha_{j} H_{j}\right)-\left[S_{2 k}\left(-\frac{i t}{r}\right)\right]^{r}\right\|=O\left(\frac{(L \Lambda t)^{2 k+1}}{r^{2 k}}\right) . \tag{6}
\end{equation*}
$$

For the $n$-qubit system described in Section 2, we have $L=O(n)$ and $\Lambda=O(1)$. Using the first-order formula to simulate that system for time $t=n$ within a fixed error $\epsilon$, it suffices to take
$r=O\left(n^{4}\right)$. The circuit for each segment has size $O(n)$, giving an overall gate complexity of $O\left(n^{5}\right)$. Similarly, the complexity of simulation using the $(2 k)$ th-order formula is $O\left(n^{3+1 / k}\right)$.

The main challenge in making these algorithms concrete is to choose an explicit $r$ that ensures some desired upper bound on the error. We derive such bounds in Appendix F. We consider four bounds, which we call the analytic bound, the minimized bound, the commutator bound, and the empirical bound. The analytic and minimized bounds involve only minor improvements over standard approaches [11], whereas the commutator bound can be significantly tighter, and the empirical bound attempts to capture the true performance.

The commutator bound, derived in Section F.2, exploits the fact that the error is smaller if many pairs of terms in the Hamiltonian commute. This bound not only improves the overall gate count of the PF algorithm, but also tightens the asymptotic complexity with respect to the system size $n$. Specifically, for the ( $2 k$ )th-order formula, this bound gives an asymptotic gate count $O\left(n^{3+2 /(2 k+1)}\right)$, improving over the above $O\left(n^{3+1 / k}\right)$.

To evaluate the commutator bound, we must compute the number of $(2 k+1)$-tuples of terms from the Hamiltonian satisfying certain commutation relations. A brute-force approach to this counting problem takes time $O\left(n^{2 k+1}\right)$, which can be prohibitive for large $n$, even for modest $k$. However, for certain Hamiltonians (including our target system), we show how to use the combinatorial structure of the bound to compute it in closed form. We evaluate the bound for the first-, second- and fourth-order product formulas, and study its asymptotic behavior for higher-order formulas. We compare the commutator bound to other PF error bounds in Section I.1.

While the analytic, minimized, and commutator bounds all provide rigorous performance guarantees, even the strongest of these - the commutator bound - is likely to be loose in practical applications. To overcome this, we consider a non-rigorous bound based on extrapolating the actual error seen in small instances using classical simulation. The details of how we compute this empirical bound are described in Section F.3.

Although we focus on improving the gate count for small instances, note that the commutator and empirical bounds actually improve the asymptotic performance of the PF algorithm as a function of system size. We discuss the nature of this improvement in Section I.1.

## C. 2 Taylor series algorithm

We now summarize the Taylor series (TS) algorithm [14]. This algorithm directly implements the (truncated) Taylor series of the evolution operator for a carefully-chosen constant time, and repeats that procedure until the entire evolution time has been simulated.

Denote the Taylor series for the evolution up to time $t$, truncated at order $K$, by

$$
\begin{equation*}
\tilde{U}(t):=\sum_{k=0}^{K} \frac{(-i t H)^{k}}{k!} . \tag{7}
\end{equation*}
$$

For sufficiently large $K$, the operator $\tilde{U}(t)$ is a good approximation of $\exp (-i H t)$. Using (2), we
can rewrite $\tilde{U}(t)$ as a linear combination of unitaries, namely

$$
\begin{align*}
\tilde{U}(t) & =\sum_{k=0}^{K} \frac{(-i t H)^{k}}{k!}  \tag{8}\\
& =\sum_{k=0}^{K} \sum_{\ell_{1}, \ldots, \ell_{k}=1}^{L} \frac{t^{k}}{k!} \alpha_{\ell_{1}} \cdots \alpha_{\ell_{k}}(-i)^{k} H_{\ell_{1}} \cdots H_{\ell_{k}}  \tag{9}\\
& =\sum_{j=0}^{\Gamma-1} \beta_{j} \tilde{V}_{j}, \tag{10}
\end{align*}
$$

for $\Gamma=\sum_{k=0}^{K} L^{k}$, where the $\tilde{V}_{j}$ are products of the form $(-i)^{k} H_{\ell_{1}} \cdots H_{\ell_{k}}$, and the $\beta_{j}$ are the corresponding coefficients such that $\beta_{j}>0$. (For notational convenience, we omit the dependence of $\beta_{j}$ on $t$.) The TS algorithm effectively implements this linear combination.

To do this, for any $t>0$ we define an isometry $\mathcal{V}(t): \mathcal{H} \rightarrow \mathbb{C}^{\Gamma} \otimes \mathcal{H}$ as follows. Let $B$ be a unitary operator on $\mathbb{C}^{\Gamma}$ satisfying

$$
\begin{equation*}
B|0\rangle=\frac{1}{\sqrt{s}} \sum_{j=0}^{\Gamma-1} \sqrt{\beta_{j}}|j\rangle \tag{11}
\end{equation*}
$$

where

$$
\begin{equation*}
s:=\sum_{j=0}^{\Gamma-1} \beta_{j}=\sum_{k=0}^{K} \frac{\left(t\left(\alpha_{1}+\cdots+\alpha_{L}\right)\right)^{k}}{k!} . \tag{12}
\end{equation*}
$$

We also define

$$
\begin{equation*}
W:=\left(B^{\dagger} \otimes I\right) \operatorname{select}(\tilde{V})(B \otimes I) \tag{13}
\end{equation*}
$$

where

$$
\begin{equation*}
\operatorname{select}(\tilde{V}):=\sum_{j=0}^{\Gamma-1}|j\rangle\langle j| \otimes \tilde{V}_{j} . \tag{14}
\end{equation*}
$$

It is easy to see that $(\langle 0| \otimes I) W(|0\rangle \otimes I) \propto \tilde{U}(t)$. More precisely, we have

$$
\begin{equation*}
W|0\rangle|\psi\rangle=\frac{1}{s}|0\rangle \tilde{U}(t)|\psi\rangle+\sqrt{1-\frac{1}{s^{2}}}|\Phi\rangle \tag{15}
\end{equation*}
$$

for $|\psi\rangle \in \mathcal{H}$ and some $|\Phi\rangle$ whose ancillary state is supported in the subspace orthogonal to $|0\rangle$. To boost the amplitude to perform the desired operation, we consider the isometry

$$
\begin{equation*}
\mathcal{V}(t):=-W R W^{\dagger} R W(|0\rangle \otimes I) \tag{16}
\end{equation*}
$$

where

$$
\begin{equation*}
R:=(I-2|0\rangle\langle 0|) \otimes I \tag{17}
\end{equation*}
$$

To ensure that $\mathcal{V}(t)$ implements evolution according to $H$ nearly deterministically, we consider evolution for time

$$
\begin{equation*}
t_{\mathrm{seg}}:=\frac{\ln 2}{\alpha_{1}+\cdots+\alpha_{L}} . \tag{18}
\end{equation*}
$$

The overall evolution is realized as a sequence of $r:=\left\lceil t / t_{\operatorname{seg}}\right\rceil$ segments, where the first $r-1$ segments each evolve the state for time $t_{\text {seg }}$ and the final segment evolves the state for time $t_{\text {rem }}:=t-(r-1) t_{\text {seg }}$. It can be shown that there is a choice of $K$ with

$$
\begin{equation*}
K=O\left(\frac{\log \frac{\left(\alpha_{1}+\cdots+\alpha_{L}\right) t_{\text {seg }}}{\epsilon}}{\log \log \frac{\left(\alpha_{1}+\cdots+\alpha_{L}\right) t_{\text {seg }}}{\epsilon}}\right) \tag{19}
\end{equation*}
$$

such that

$$
\begin{equation*}
\|(\langle 0| \otimes I) \mathcal{V}\left(t_{\mathrm{seg}}\right)-\exp \left(-i t_{\mathrm{seg}} H\right) \|=O(\epsilon / r) . \tag{20}
\end{equation*}
$$

The evolution for the remaining time $t_{\text {rem }}$ can be performed by rotating an ancilla qubit to artificially increase the duration of the segment. Specifically, provided $s<2$, we can introduce an ancilla register in the state $|0\rangle$ and apply the rotation

$$
\begin{equation*}
|0\rangle \mapsto \frac{s}{2}|0\rangle+\sqrt{1-\frac{s^{2}}{4}}|1\rangle . \tag{21}
\end{equation*}
$$

Together with the the unitary operator $W$, this implements the transformation

$$
\begin{equation*}
|0\rangle|\psi\rangle \mapsto \frac{1}{2}|00\rangle \tilde{U}(t)|\psi\rangle+\frac{\sqrt{3}}{2}\left|\Phi^{\prime}\right\rangle \tag{22}
\end{equation*}
$$

for some normalized state $\left|\Phi^{\prime}\right\rangle$ with $(\langle 00| \otimes I)\left|\Phi^{\prime}\right\rangle=0$. Then we can proceed as before, but with $s=2$. Indeed, we also perform a similar rotation for the initial $r-1$ segments to ensure that they have $s=2$ instead of a slightly smaller value. By an abuse of notation, we incorporate this rotation into the definition of $\mathcal{V}$, so that $\mathcal{V}\left(t_{\text {seg }}\right)$ and $\mathcal{V}\left(t_{\text {rem }}\right)$ are the corresponding evolution operators for the first $r-1$ segments and the final segment.

The asymptotic gate complexity of this simulation algorithm is [14]

$$
\begin{equation*}
O\left(T L(n+\log L) \frac{\log (T / \epsilon)}{\log \log (T / \epsilon)}\right) \tag{23}
\end{equation*}
$$

where $T=\left(\alpha_{1}+\cdots+\alpha_{L}\right) t$. For the Hamiltonian studied in this paper, we have $T=O\left(n^{2}\right)$ and $L=O(n)$, which gives a bound of $O\left(n^{4} \frac{\log n}{\log \log n}\right)$. However, the analysis in [14] assumes only that each term in the Hamiltonian is a tensor product of Pauli gates, possibly acting nontrivially on all of the qubits. Since our Hamiltonian is 2-local, we have a tighter bound of $O\left(n^{3} \frac{\log ^{2} n}{\log \log n}\right)$ for the asymptotic gate count, as indicated in Table 1.

To concretely implement the TS algorithm, we must replace the asymptotic statements above with an explicit error analysis. We present the details of such an analysis in Appendix G. In particular, Section G. 1 shows how to choose $K$ to ensure that the overall error is at most some allowed $\epsilon$. In addition, unlike the PF algorithm, the TS algorithm requires a measurement on the ancilla register and succeeds probabilistically. We discuss in Section G. 2 how this fact can be taken into account to make a fair comparison between the PF and TS algorithms.

Another crucial aspect is the implementation of the select $(\tilde{V})$ gate. In Section G.3, we discuss how we encode the control register for this operation. Then, in Section G.4, we present a novel method to implement select $(V)$ gates by walking on a binary tree, improving the gate complexity to $O(\Gamma)$ from the naive complexity of $O(\Gamma \log \Gamma)$.

It is also natural to ask whether an empirical error bound could be established for the TS algorithm. Unfortunately, the number of ancilla qubits used by the algorithm (as shown in Figure 2) makes direct classical simulation infeasible even for very small sizes. An alternative is to only give an empirical bound on the remainder of the Taylor series. However, as we discuss in Section G.1, such a bound does not give a significant improvement.

## C. 3 Quantum signal processing algorithm

Now we summarize the quantum signal processing (QSP) algorithm of Low and Chuang [51, 52]. Again consider a Hamiltonian of the form (2). We have

$$
\begin{equation*}
H=\alpha(\langle G| \otimes I) \operatorname{select}(H)(|G\rangle \otimes I) \tag{24}
\end{equation*}
$$

where the first register holds an $L$-dimensional ancilla,

$$
\begin{equation*}
\operatorname{select}(H):=\sum_{\ell=1}^{L}|\ell\rangle\langle\ell| \otimes H_{\ell} \tag{25}
\end{equation*}
$$

(similarly to (14)), and

$$
\begin{equation*}
|G\rangle:=\frac{1}{\sqrt{\alpha}} \sum_{\ell=1}^{L} \sqrt{\alpha_{\ell}}|\ell\rangle, \quad \alpha:=\sum_{\ell=1}^{L} \alpha_{\ell} . \tag{26}
\end{equation*}
$$

Low and Chuang's concept of qubitization [51] relates the spectral decompositions of $H / \alpha$ and

$$
\begin{equation*}
-i Q:=-i((2|G\rangle\langle G|-I) \otimes I) \text { select }(H) \tag{27}
\end{equation*}
$$

Specifically, let $H / \alpha=\sum_{\lambda} \lambda|\lambda\rangle\langle\lambda|$ be a spectral decomposition of $H / \alpha$, where the sum runs over all eigenvalues of $H / \alpha$. By the triangle inequality, $\|H\| \leq \alpha$, i.e., $|\lambda| \leq 1$. For each eigenvalue $\lambda \in$ $(-1,1)$ of $H / \alpha$, the qubitization theorem (Theorem 2 of [51]) asserts that $-i Q$ has two corresponding eigenvalues

$$
\begin{equation*}
\lambda_{ \pm}=\mp \sqrt{1-\lambda^{2}}-i \lambda=\mp e^{ \pm i \arcsin \lambda} \tag{28}
\end{equation*}
$$

with eigenvectors $\left|\lambda_{ \pm}\right\rangle=\left(\left|G_{\lambda}\right\rangle \pm i\left|G_{\lambda}^{\perp}\right\rangle\right) / \sqrt{2}$, where

$$
\begin{equation*}
\left|G_{\lambda}\right\rangle:=|G\rangle \otimes|\lambda\rangle \tag{29}
\end{equation*}
$$

and

$$
\begin{equation*}
\left|G_{\lambda}^{\perp}\right\rangle:=\frac{\lambda\left|G_{\lambda}\right\rangle-\operatorname{select}(H)\left|G_{\lambda}\right\rangle}{\sqrt{1-\lambda^{2}}} \tag{30}
\end{equation*}
$$

(Eigenvalues $\lambda_{ \pm}= \pm 1$ correspond to degenerate cases that can be analyzed separately.)
The signal processing algorithm applies a sequence of operations called phased iterates. We introduce an additional ancilla qubit and define the operator

$$
\begin{equation*}
V_{\phi}:=\left(e^{-i \phi \sigma^{z} / 2} \otimes I\right)(|+\rangle\langle+| \otimes I+|-\rangle\langle-| \otimes(-i Q))\left(e^{i \phi \sigma^{z} / 2} \otimes I\right) \tag{31}
\end{equation*}
$$

for any $\phi \in \mathbb{R}$. Let $-i Q=\sum_{\nu} e^{i \theta_{\nu}}|\nu\rangle\langle\nu|$ be a spectral decomposition of $-i Q$, where the sum runs over $\nu$ labeling all eigenvectors of $-i Q$. As described above, each eigenvalue $\lambda \in(-1,1)$ of $H / \alpha$ corresponds to two eigenvalues $e^{i \theta_{\lambda_{ \pm}}}$of $-i Q$, where $\theta_{\lambda_{+}}=\arcsin (\lambda)+\pi$ and $\theta_{\lambda_{-}}=-\arcsin (\lambda)$. Eigenvalues $\pm 1$ of $-i Q$ correspond to degenerate cases that can be handled separately. The remaining eigenspaces cannot be reached during any execution of the quantum signal processing algorithm, so we can neglect them. Then one can show that

$$
\begin{equation*}
V_{\phi}=\sum_{\nu} e^{i \theta_{\nu} / 2} R_{\phi}\left(\theta_{\nu}\right) \otimes|\nu\rangle\langle\nu| \tag{32}
\end{equation*}
$$

where

$$
\begin{equation*}
R_{\phi}(\theta):=e^{-i \theta \sigma^{\phi} / 2}, \quad \sigma^{\phi}:=\cos (\phi) \sigma^{x}+\sin (\phi) \sigma^{y} . \tag{33}
\end{equation*}
$$

Thus each eigenvalue $e^{i \theta_{\nu}}$ of $-i Q$ is manifested in $V_{\phi}$ as an $\mathrm{SU}(2)$ operator $R_{\phi}\left(\theta_{\nu}\right)$ acting on the ancilla qubit.

For any positive even integer $M$, composing gates with the same rotation amplitude $\theta$ but with varying phases $\phi_{1}, \ldots, \phi_{M}$ yields

$$
\begin{equation*}
R_{\phi_{M}}(\theta) \cdots R_{\phi_{1}}(\theta)=A\left(\cos \frac{\theta}{2}\right) I+i B\left(\cos \frac{\theta}{2}\right) \sigma^{z}+i \cos \frac{\theta}{2} C\left(\sin \frac{\theta}{2}\right) \sigma^{x}+i \cos \frac{\theta}{2} D\left(\sin \frac{\theta}{2}\right) \sigma^{y} \tag{34}
\end{equation*}
$$

for polynomials $A, B, C, D$ of degree at most $M$. In the QSP algorithm, only the polynomials $A$ and $C$ are used. This component can be extracted by preparing the ancilla qubit in the state
 unwanted factor $e^{i \theta_{\nu} / 2}$ may be canceled by alternating between $V_{\phi}$ and $V_{\phi+\pi}^{\dagger}$, giving

$$
\begin{equation*}
V:=V_{\phi_{M}+\pi}^{\dagger} V_{\phi_{M-1}} \cdots V_{\phi_{2}+\pi}^{\dagger} V_{\phi_{1}} . \tag{35}
\end{equation*}
$$

To perform Hamiltonian simulation by qubitization, we implement a function of $\theta$ that converts the eigenvalue $e^{i \theta_{\lambda_{ \pm}}}$of $-i Q$ to the desired phase $e^{-i \lambda t}$, namely the Jacobi-Anger expansion

$$
\begin{equation*}
e^{i \sin (\theta) t}=\sum_{k=-\infty}^{\infty} J_{k}(t) e^{i k \theta} \tag{36}
\end{equation*}
$$

To do this with a polynomial of degree $M$, we truncate the expansion at order $q:=\frac{M}{2}+1$, giving an approximation with error at most [15]

$$
\begin{equation*}
2 \sum_{k=q}^{\infty}\left|J_{k}(t)\right| \leq \frac{4 t^{q}}{2^{q} q!} \tag{37}
\end{equation*}
$$

The angles $\phi_{1}, \ldots, \phi_{M}$ that realize this expansion can be computed by an efficient classical procedure (see Lemmas 1 and 3 of [53]).

To simulate evolution of an initial state $|\psi\rangle$, the QSP algorithm applies $V$ to the state $|+\rangle \otimes$ $|G\rangle \otimes|\psi\rangle$ and postselects the ancilla register of the output on the state $|+\rangle \otimes|G\rangle$. This procedure simulates the desired evolution with error at most

$$
\begin{equation*}
8 \frac{4(\alpha t)^{q}}{2^{q} q!} \leq \epsilon \tag{38}
\end{equation*}
$$

To achieve simulation for time $t$ and error $\epsilon$, the QSP algorithm uses

$$
\begin{equation*}
M=O\left(\alpha t+\frac{\log (1 / \epsilon)}{\log \log (1 / \epsilon)}\right) \tag{39}
\end{equation*}
$$

phased iterates [52]. For each phased iterate, the dominant part is the select $(H)$ subroutine, which is straightforward to implement with $O(n \log n)$ elementary gates. Overall, we see that the asymptotic gate count in terms of the system size is $O\left(n^{3} \log n\right)$. (Note that by using our improved select (•) implementation described in Section G.4, the asymptotic complexity is reduced to $O\left(n^{3}\right)$.)

Operationally, post-selection of the ancilla is achieved by measurement. If the outcome is
 we compute the success probability of the QSP algorithm and discuss how it can be fairly compared with the deterministic PF algorithm.

When implementing the QSP algorithm, we eliminate unnecessary gates wherever possible to reduce the gate count. In particular, note that in place of (31), Low and Chuang define the phased iterate as

$$
\begin{equation*}
V_{\phi}^{\prime}:=\left(e^{-i \phi \sigma^{z} / 2} \otimes I\right)\left(|+\rangle\langle+| \otimes I+|-\rangle\langle-| \otimes\left(Z_{-\pi / 2}(-i Q) Z_{\pi / 2}\right)\right)\left(e^{i \phi \sigma^{z} / 2} \otimes I\right), \tag{40}
\end{equation*}
$$

where $Z_{\varphi}:=\left(1+e^{-i \varphi}\right)|G\rangle\langle G|-I$ is a partial reflection about $|G\rangle$. Our modified definition (31) saves two partial reflections in each phased iterate - saving $O\left(n^{2} \log n\right)$ gates overall-but has the same behavior. This and other optimizations of the implementation are detailed in Section H.2.

The QSP algorithm requires more substantial classical preprocessing than the PF and TS approaches. Computing the angles $\phi_{1}, \ldots, \phi_{M}$ requires finding the roots of a polynomial of degree $2 M$, and these roots must be determined to high precision. Thus we were unable to compute the parameters of the algorithm explicitly except in very small instances.

As discussed in Section 3.3, we address this issue by considering a segmented version of the algorithm. We discuss this approach further in Section H.3. Here we briefly consider how the segmented implementation impacts the asymptotic performance of the QSP algorithm.

Suppose we fix a positive even integer $M$, the maximal number of phased iterates for which the angles $\phi_{1}, \ldots, \phi_{M}$ can practically be computed. As shown in Section H.3, it suffices to use $r=O\left(t\left(\frac{t}{\epsilon}\right)^{2 / M}\right)$ segments to ensure overall error at most $\epsilon$. In the instance of Hamiltonian simulation considered in this paper, with $t=n$ and $\alpha=O(n)$, we have $r=O\left(n^{2+4 / M}\right)$ segments. Within each segment, the number of phased iterates is $M$, which is independent of the system size. The circuit size of each phased iterate is $O(n)$ using the improved select $(V)$ implementation described in Section G.4. Thus the segmented algorithm has gate complexity $O\left(n^{3+4 / M}\right)$. In our implementation of the segmented algorithm, we use $M=28$, so the exponent is about 3.14.

We also consider empirical bounds for the QSP algorithm. Specifically, we find an improved empirical estimate of the truncation error of the Jacobi-Anger expansion. This partial empirical bound leads to a small reduction in the gate count, as discussed further in Section H.4. As with the TS algorithm, the need for ancilla qubits in the QSP algorithm makes it difficult to establish a comprehensive empirical bound by performing a full simulation on a classical computer. Fixing the target error $\epsilon=10^{-3}$, we investigate the empirical performance of QSP for systems of size 5 to 9. Our preliminary data suggest that the gate count is not significantly improved even with such an empirical bound, so we do not consider such a bound in our study.

## D System-size dependence for other simulation algorithms

In this appendix, we discuss the asymptotic gate complexity of the Hamiltonian simulation algorithms we chose not to implement.

For a local Hamiltonian, the gate complexity of the algorithm based on fractional queries is [13]

$$
\begin{equation*}
O\left(\tilde{\tau} \frac{\log (\tilde{\tau} / \epsilon)}{\log \log (\tilde{\tau} / \epsilon)}(\log (\tilde{\tau} / \epsilon)+n)\right) \tag{41}
\end{equation*}
$$

where $\tilde{\tau}=L\|H\|_{\max } t$, with $\|H\|_{\max }$ denoting the largest entry of $H$ in absolute value (note that this expression slightly tightens the one given in the main statement of Theorem 1.1 of [13]; see the end of the proof of that theorem for details). We have $L=\Theta(n)$ and $\|H\|_{\max }=\Theta(n)$, so $\tilde{\tau}=O\left(n^{3}\right)$, and the gate complexity is $O\left(n^{4} \frac{\log n}{\log \log n}\right)$.

Similarly, the gate complexity of the algorithm based on a linear combination of quantum walk steps is [15, Theorem 1]

$$
\begin{equation*}
O\left(\tau\left(n+\log ^{5 / 2}(\tau / \epsilon)\right) \frac{\log (\tau / \epsilon)}{\log \log (\tau / \epsilon)}\right) \tag{42}
\end{equation*}
$$

where $\tau=d\|H\|_{\max } t$ with $d$ denoting the sparsity of $H$. Since $d=\Theta(n)$, we find a gate complexity of $O\left(n^{4} \frac{\log n}{\log \log n}\right)$.

We now turn to the simulation algorithm based on quantum walk [12, Theorem 1]. Previous work on this approach has not focused on the gate complexity, so we evaluate it here. The algorithm
proceeds as follows. First we perform phase estimation on the quantum walk with one bit of precision. We then apply $O\left(d\|H\|_{\max } t\right)$ steps of a lazy quantum walk. Finally, we invoke phase estimation again with $O\left(d\|H\|_{\max } t\right)$ applications of the walk operator and use this estimate to correct the phase, further improving the accuracy.

The above procedure uses $O\left(d\|H\|_{\max } t\right)=O\left(n^{3}\right)$ quantum walk steps, which dominates the cost of the Fourier transform in the phase estimation procedure (since the estimated phase has $\log \left(d\|H\|_{\max } t\right)=O(\log n)$ bits, the Fourier transform has complexity $\left.O(\log n \log \log n)\right)$ and the coherent computation of the sine function to correct the phase (with complexity $O(M(\log n) \log \log n)=$ poly $(\log n)$, where $M(k)$ is the complexity of multiplying $k$-bit numbers [17]). Thus, to understand the asymptotic gate complexity of the algorithm, it suffices to study the gate complexity of performing a quantum walk step.

The quantum walk operator is the product of a swap (with complexity $O(n)$ ) and the reflection

$$
\begin{equation*}
\sum_{x=1}^{2^{n}}|x\rangle\langle x| \otimes\left(2\left|\phi_{x}\right\rangle\left\langle\phi_{x}\right|-I\right), \tag{43}
\end{equation*}
$$

where

$$
\begin{equation*}
\left|\phi_{x}\right\rangle:=\frac{1}{\sqrt{2^{n}}} \sum_{y=1}^{2^{n}}|y\rangle\left[\sqrt{\frac{H_{x y}^{*}}{\mathcal{X}}}|0\rangle+\sqrt{1-\frac{\left|H_{x y}\right|}{\mathcal{X}}}|1\rangle\right] \tag{44}
\end{equation*}
$$

with

$$
\begin{equation*}
\mathcal{X}:=\frac{1}{d t} \max \left\{\lceil\|H\| t / \sqrt{\epsilon}\rceil,\left\lceil\|H\|_{\max } d t\right\rceil\right\} \tag{45}
\end{equation*}
$$

To implement the walk operator, it suffices to give procedures for preparing $\left|\phi_{x}\right\rangle$ and reflecting about $|0\rangle$. The reflection operator can be implemented by performing an $X$ gate on each qubit, applying a controlled- $Z$ gate with $n-1$ control qubits, and again performing an $X$ gate on each qubit, for a total cost of $O(n)$. Thus it remains to understand the cost of preparing $\left|\phi_{x}\right\rangle$.

The algorithm uses oracles $O_{F}$ and $O_{H}$ acting as

$$
\begin{equation*}
O_{F}|x, \ell\rangle=|x, f(x, \ell)\rangle \tag{46}
\end{equation*}
$$

and

$$
\begin{equation*}
O_{H}|x, y, z\rangle=\left|x, y, z \oplus H_{x y}\right\rangle, \tag{47}
\end{equation*}
$$

where $f(x, \ell)$ is the column index of the $\ell$ th nonzero element in row $x$. Reference [12] observes that $\left|\phi_{x}\right\rangle$ can be prepared with cost

$$
\begin{equation*}
O\left(n+\operatorname{cost}\left(O_{F}\right)+\operatorname{cost}\left(O_{H}\right)\right) . \tag{48}
\end{equation*}
$$

Therefore, it suffices to separately understand the cost of implementing $O_{F}$ and $O_{H}$ for our Hamiltonian.

Recall that the Hamiltonian takes the form

$$
\begin{equation*}
H=\sum_{j=1}^{n}\left(\vec{\sigma}_{j} \cdot \vec{\sigma}_{j+1}+h_{j} \sigma_{j}^{z}\right) \tag{49}
\end{equation*}
$$

For any $j \in\{1, \ldots, n\}$, the term $\vec{\sigma}_{j} \cdot \vec{\sigma}_{j+1}$ acts on the $j$ th and $(j+1)$ st qubits according to the matrix

$$
\left[\begin{array}{cccc}
1 & 0 & 0 & 0  \tag{50}\\
0 & -1 & 2 & 0 \\
0 & 2 & -1 & 0 \\
0 & 0 & 0 & 1
\end{array}\right] .
$$

Thus the nonzero elements of row $x$ of $H$, aside from the diagonal, correspond to the ways of swapping two adjacent bits in the binary representation of $x$. There are at most $n+1$ such nonzero elements, so we can suppose that $\ell \in\{1, \ldots, n+1\}$ in (46), where (say) $\ell=n+1$ corresponds to the diagonal element.

To implement $O_{F}$, we begin by making a copy of the first register using $O(n)$ CNOT gates, performing

$$
\begin{equation*}
|x, \ell\rangle \mapsto|x, x, \ell\rangle . \tag{51}
\end{equation*}
$$

Then, conditioned on the value of $\ell \in\{1, \ldots, n\}$, we swap the $\ell$ th and $(\ell+1)$ st bits of the second register, giving

$$
\begin{equation*}
|x, x, \ell\rangle \mapsto|x, f(x, \ell), \ell\rangle . \tag{52}
\end{equation*}
$$

This can be done with complexity $O(n \log n)$, since it is a select $(\cdot)$ gate where the selection register has $n+1$ possible states and the target gates act only on two qubits. To uncompute the third register, note that there is a classical algorithm with running time $O(n)$ that computes $\ell$ from $x$ and $f(x, \ell)$ : we just compare the two strings and detect which pair of bits are swapped. This procedure can be made reversible by standard techniques. Therefore, we can erase $|\ell\rangle$ and thereby implement $O_{F}$ with gate complexity $O(n \log n)$.

To implement $O_{H}$, we use a classical algorithm that, given the row and column indices $(x, y)$, outputs the corresponding matrix element $H_{x y}$. The $(x, x)$ diagonal element is easy to compute from the binary representation of $x$ : every adjacent pair of bits gives a contribution of +1 if the bits agree and -1 if they disagree, and the $j$ th bit gives an additional contribution of $h_{j}$ if the bit is 0 and $-h_{j}$ if the bit is 1 . For $x \neq y$, the off-diagonal element $H_{x y}$ is 2 if $x$ and $y$ differ by swapping some adjacent pair of distinct bits, and is 0 otherwise. Since all these calculations can be performed in time $O(n)$, the complexity of implementing $O_{H}$ is $O(n)$.

Altogether, we see that the gate complexity of the quantum walk simulation is $O\left(n^{4} \log n\right)$, as shown in Table 1.

In fact, our improved implementation of select(•) gates in Lemma G. 7 shows that $O_{F}$ can be implemented in time $O(n)$. With this improvement, the quantum walk simulation can be performed slightly faster, in time $O\left(n^{4}\right)$.

## E Circuit synthesis and optimization

We implemented quantum simulation algorithms in the Quipper programming language [32]. Quipper is a circuit description language equipped with many high-level circuit combinators (e.g., circuit iteration and circuit reversal) that allow for a concise specification of complex quantum circuits. Quipper also supports a hierarchical circuit structure that allows for the efficient manipulation of the very large quantum circuits considered here. We made full use of these features. The Quipper source code for our implementations, together with sample output circuits and optimized versions thereof, are available in a public repository [26].

## E. 1 Overview of algorithm implementations

## E.1.1 Product formulas

We implement simulation algorithms using product formulas of order 1, 2, 4, 6, and 8. For each order, we implement various time-slicing methods as discussed in Section C. 1 and Appendix F. Specifically, we implement the analytic, minimized, and empirical bounds for all orders. For first, second-, and fourth-order algorithms, we also implement the commutator bound. Overall, we consider 18 different types of product formula algorithm.

The product formula algorithms are the simplest of the simulation algorithms we consider. To simulate the evolution for a short time, these algorithms simply exponentiate each term of the Hamiltonian (which in our case is always proportional to a 1- or 2-qubit Pauli operator) in a carefully-chosen sequence. Simulation for a longer time is then obtained by iterating this sequence. The duration of the individual evolutions and the total number of repetitions is determined using an appropriate error bound. The iteration is easily realized using Quipper's built-in iteration combinator.

## E.1.2 Taylor series

The TS algorithm is more involved than the PF algorithm, using more complicated subroutines. Our implementation is based on a mixed unary-binary encoding of the control register, as explained in Section G.3. The implementation consists of several subroutines, including two distinct state preparation subroutines, the select $(V)$ procedure discussed in Section G.4, and the reflection about $|0\rangle$. The first state preparation subroutine maps $|0\rangle$ into the state $\sum_{l=1}^{L} \sqrt{\alpha_{l}}|l\rangle$, where $\left\{\alpha_{l}\right\}_{l=1}^{L}$ are the coefficients of terms in the target Hamiltonian. This is accomplished through the generic state preparation algorithm described in [69]. The second state preparation subroutine generates a state proportional to $\sum_{k=0}^{K} \sqrt{t^{k} / k!}\left|1^{k} 0^{K-k}\right\rangle$ starting with the basis state $|0\rangle$, for which the generic method [69] is suboptimal. Instead, we use the state preparation procedure described in [14], applying a rotation on the first qubit, followed by rotations on qubits $k=2$ to $K$ controlled by the qubit $k-1$. We develop our own implementation of the $\operatorname{select}(V)$ operation, as described in Section G.4. The reflection about $|0\rangle$ is implemented as a multiply-controlled $Z$ gate, following the construction in [55] that uses ancillas in the state $|0\rangle$.

## E.1.3 Quantum signal processing

We implement two versions of the QSP algorithm. First, we consider the original version of the algorithm as described in Section C.3. In choosing the parameters for this implementation, we use the empirical estimate of the remainder of the Jacobi-Anger expansion described in Section H.4. Unfortunately, our implementation of the classical computation of the rotation angles [53] is only able to handle very small system sizes. Instead, to generate gate counts, we use randomly-selected angles as placeholder values. Thus the circuits constructed by our implementation have the correct structure (and in particular, give the correct gate count over the Clifford $+R_{z}$ basis, and a good approximation over the Clifford $+T$ basis), but do not actually implement the desired unitary.

We also implement a segmented version of the QSP algorithm as discussed near the end of Section C. 3 and detailed in Section H.3. In this case, we invoke a rigorous error bound, and we are able to correctly compute all parameters of the algorithm. However, the rotation angle computation is involved and uses high-precision arithmetic. For this reason, we compute these parameters off-line using Mathematica and store the results in a look-up table that is accessed by the Quipper program to construct the quantum circuit. Thus our Quipper implementation can only produce circuits for the system sizes for which these values have been precomputed. The Mathematica scripts are available as part of our implementation [26] so that the interested reader can compute the required parameters and add them to the look-up tables if additional system sizes are of interest.

## E. 2 Gate counts and synthesis

We express all algorithms using Clifford gates (including CNOT) and single-qubit z-rotations by arbitrary angles. We use Quipper's standard gate-counting feature to produce the first set of resource estimates (in the Clifford $+R_{z}$ basis). To obtain our second set of estimates (in the Clifford $+T$ basis),
we must approximate $z$-rotations by Clifford $+T$ circuits. These approximations are obtained using the optimal algorithm of Ross and Selinger [66] (which builds upon the exact synthesis algorithm of Kliuchnikov, Maslov, and Mosca [45]).

Note that a better approach to decomposing $R_{z}$ gates into Clifford $+T$ circuits might be to rely on the repeat-until-success (RUS) strategy of [16]. This reduces the $T$-count by a factor of about 2.5 on average, at the cost of using additional resources in the form of measurements, classical feedback, and additional cnot gates. One may combine RUS decomposition with Campbell's unitary mixing approach [20] to further reduce the $T$-count by a factor of about 2 . Finally, rather than using only the $T=R_{z}(\pi / 4)$ gate, one may also distill $R_{z}(\pi / 6)$ with comparable or better efficiency as the $T$ gate [18]. Using $R_{z}(\pi / 6)$ along with the $T$ gate in the RUS circuits [16] together with unitary mixing [20], we expect an improvement in the resource count for a fault-tolerant implementation by a factor of 5 or more. Alternatively, one might employ the gate distillation techniques of [29], which could also significantly reduce the cost of implementing fault-tolerant gates. We leave a detailed investigation of such possible improvements as a subject for future work.

When producing Clifford $+T$ gate counts for the non-segmented QSP algorithm, we synthesize rotations that are incorrect, since we are unable to compute the exact rotation angles. Since the number of Clifford $+T$ gates required to synthesize a given rotation depends on its angle, the resulting gate counts are not, strictly speaking, precise. Nevertheless, we expect the produced gate counts to accurately represent the true counts. Since the rotation angles are randomly chosen, their cost is that of a typical angle (roughly $3 \log (1 / \tau)$ where $\tau$ is the approximation precision). Furthermore, because we synthesize many rotations, occurrences of over/underestimation of the specific gate counts average out over the entire circuit by the law of large numbers. The cost of the approximation therefore depends primarily on the precision - which, in turn, depends on the number of rotations in the overall circuit-but not on the individual rotation angles.

To approximate a given Hamiltonian $e^{-i H t}$ to precision $\epsilon$ with a Clifford $+T$ circuit, we divide the allowed error $\epsilon$ between gate synthesis and simulation algorithm errors. In our implementation, we allocate $\epsilon / 2$ to the simulation error and $\epsilon / 2$ to the gate synthesis error. The latter is then divided by the total number of $z$-rotations appearing in the circuit. While anecdotal evidence suggests that such an even division of $\epsilon$ between simulation and synthesis is a good choice, we leave it as an avenue for future research to determine whether a more subtle partition should be preferred.

## E. 3 Optimization

We also employ a circuit optimizer [57] that uses a variety of techniques to reduce gate counts in an automated way. The optimizer uses various circuit equivalences to induce quantum gate commutations, mergers, and cancellations.

Many of the subcircuits found in the TS and QSP algorithms are based on Toffoli gates. In our unoptimized circuits, we simply replace each Toffoli gate with an optimal Clifford $+T$ implementation [2]. However, there are many ways to write the Toffoli gate as a Clifford $+T$ circuit (e.g., some gates in an implementation may commute, controls can be interchanged, and the circuit can be reversed and/or complex conjugated since it is real and self-inverse). Carefully choosing when to use which decomposition can enable additional gate cancellations. For this reason, we outsource Toffoli gate decomposition to the optimizer.

Once the circuits are expressed over the Clifford $+R_{z}$ basis, our optimizer first performs a sequence of rewrites to reduce the number of Hadamard gates. The resulting circuit is more amenable to further optimization since it typically contains larger chunks that can be expressed using so-called phase polynomials $[2,57]$. Then we identify pairs of gates that can be canceled or merged (allowing the gates in each pair to be separated by a subcircuit through which one of the gates commutes,
according to a fixed set of commutation rules). Finally, we use the phase polynomial representation over $\left\{\right.$ CNOT, $\left.R_{z}\right\}$ to further lower the $R_{z}$ count of the circuit. This representation can be used to identify $R_{z}$ gates that are applied to the same linear Boolean function of the input, and thereby merge them, even if they originally correspond to distant gates. Merging such rotations can enable additional gate commutations, leading to more optimizations. We repeat the entire sequence of optimization procedures until no more gate count reduction is achieved.

Our optimizer comes with "light" and "heavy" options [57]. In general, the heavy version finds more reductions, but is more computationally intensive. We used the light version of the optimizer to obtain the results reported in this paper since we do not expect the heavy version to change our conclusions qualitatively.

## E. 4 Correctness

We carry out a number of tests to verify correctness of our circuit-level implementations. We simulate entire PF circuits for systems with up to 15 qubits and check for proximity to the ideal evolution operator by evaluating the spectral norm distance. We also simulate the entire segmented QSP circuit for a system of size 5 and compare the outcome of the circuit to the ideal operator in the same sense. Unfortunately, the number of ancillas makes full simulation for larger instances of the TS algorithm prohibitively expensive. Thus, in lieu of full simulation, we test individual subroutines for correctness: we verify the reflection and state preparation subroutines for systems of up to 15 qubits, and test various instances of the select $(V)$ circuit, including its controlled versions, for systems of size 5. Finally, we test the output of the optimizer by simulating and comparing circuits before and after optimization [57].

## F Product formula implementation details

The main issue involved in implementing product formula (PF) algorithms-as introduced in Section C.1-is to establish error bounds that determine how finely to split the evolution so that some target error is achieved. In this section, we present explicit error analysis for PF algorithms that yields effective methods for computing values of $r$ that ensure the error is at most $\epsilon$.

First, we establish error bounds that provide a proven guarantee that the simulated evolution is $\epsilon$-close to the ideal evolution. We present three such bounds, which we call the analytic, minimized, and commutator bounds. The analytic and minimized bounds, presented in Section F.1, follow essentially the same strategy as in previous work [11]. The commutator bound, presented in Section F.2, is substantially more involved, using the structure of commutators between terms in the Hamiltonian to improve the result.

Second, in Section F.3, we present empirical error bounds, which are obtained by extrapolating numerical data. Since they ostensibly describe the true performance of PF algorithms, the empirical bounds result in smaller gate counts than the rigorous bounds. However, while the empirical bounds are plausible, they do not provide a guarantee about the actual distance between the simulated evolution and the ideal evolution, unlike with rigorous bounds.

Throughout, we quantify the simulation error with respect to the spectral norm. Note that for a unitary process, the diamond norm distance is at most twice the spectral norm distance [15, Lemma 7].

## F. 1 Analytic and minimized bounds

Reference [11] gives an explicit error bound for Suzuki product formulas. Using similar techniques, we can also give an explicit bound for the first-order case. We present these bounds here, tightening the analysis wherever possible.

First, recall some useful properties of Taylor expansions. For any $k \in \mathbb{N}$ and any analytic function $f: \mathbb{C} \rightarrow \mathbb{C}$ with $f(x)=\sum_{j=0}^{\infty} a_{j} x^{j}$, let $\mathcal{R}_{k}(f):=\sum_{j=k+1}^{\infty} a_{j} x^{j}$ denote the remainder of the Taylor series expansion of $f$ up to order $k$.

Lemma F.1. If $\lambda \in \mathbb{C}$ and $H_{1}, \ldots, H_{L}$ are Hermitian, then

$$
\begin{equation*}
\left\|\mathcal{R}_{k}\left(\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right)\right\| \leq \mathcal{R}_{k}\left(\exp \left(\sum_{j=1}^{L}|\lambda| \cdot\left\|H_{j}\right\|\right)\right) . \tag{53}
\end{equation*}
$$

Proof. We have

$$
\begin{equation*}
\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)=\prod_{j=1}^{L} \sum_{r_{j}=0}^{\infty} \lambda^{r_{j}} \frac{H_{j}^{r_{j}}}{r_{j}!}=\sum_{r_{1}, \ldots, r_{L}=0}^{\infty} \prod_{j=1}^{L} \lambda^{r_{j}} \frac{H_{j}^{r_{j}}}{r_{j}!}, \tag{54}
\end{equation*}
$$

so

$$
\begin{equation*}
\mathcal{R}_{k}\left(\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right)=\sum_{\substack{r_{1}, \ldots, r_{L}=0 \\ \sum_{l}, r_{l} \geq k+1}}^{\infty} \prod_{j=1}^{L} \lambda^{r_{j}} \frac{H_{j}^{r_{j}}}{r_{j}!} . \tag{55}
\end{equation*}
$$

Using the triangle inequality and submultiplicativity of the norm, we find

$$
\begin{equation*}
\left\|\mathcal{R}_{k}\left(\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right)\right\|=\left\|\sum_{\|}^{\infty} \prod_{\substack{r_{1}, \ldots, r_{L}=0 \\ \sum_{l} r_{l} \geq k+1}}^{L} \lambda^{r_{j}} \frac{H_{j}^{r_{j}}}{r_{j}!}\right\| \leq \sum_{\substack{r_{1}, \ldots, r_{L}=0 \\ \sum_{l} r_{l} \geq k+1}}^{\infty} \prod_{j=1}^{L}|\lambda|^{r_{j}} \frac{\left\|H_{j}\right\|^{r_{j}}}{r_{j}!} . \tag{56}
\end{equation*}
$$

Finally, similarly as in (54) and (55), we have

$$
\begin{equation*}
\sum_{\substack{r_{1}, \ldots, r_{L}=0 \\ \sum_{l} r_{l} \geq k+1}}^{\infty} \prod_{j=1}^{L}|\lambda|^{r_{j}} \frac{\left\|H_{j}\right\|^{r_{j}}}{r_{j}!}=\mathcal{R}_{k}\left(\prod_{j=1}^{L} \exp \left(|\lambda| \cdot\left\|H_{j}\right\|\right)\right)=\mathcal{R}_{k}\left(\exp \left(\sum_{j=1}^{L}|\lambda| \cdot\left\|H_{j}\right\|\right)\right) \tag{57}
\end{equation*}
$$

which completes the proof.
Lemma F.2. If $\lambda \in \mathbb{C}$, then $\left|\mathcal{R}_{k}(\exp (\lambda))\right| \leq \frac{|\lambda|^{k+1}}{(k+1)!} \exp (|\lambda|)$.
Proof. Using the Taylor expansion of the exponential function, we have

$$
\begin{align*}
\left|\mathcal{R}_{k}(\exp (\lambda))\right| & =\left|\sum_{l=k+1}^{\infty} \frac{\lambda^{l}}{l!}\right| \leq \sum_{l=k+1}^{\infty} \frac{|\lambda|^{l}}{l!}=|\lambda|^{k+1} \sum_{l=k+1}^{\infty} \frac{|\lambda|^{l-(k+1)}}{l!}  \tag{58}\\
& =|\lambda|^{k+1} \sum_{l=0}^{\infty} \frac{|\lambda|^{l}}{(l+(k+1))!} \leq|\lambda|^{k+1} \sum_{l=0}^{\infty} \frac{|\lambda|^{l}}{l!} \frac{1}{(k+1)!}  \tag{59}\\
& =\frac{|\lambda|^{k+1}}{(k+1)!} \exp (|\lambda|), \tag{60}
\end{align*}
$$

which completes the proof.

Now we give an explicit error bound for the first-order product formula.
Proposition F.3. Let $r \in \mathbb{N}$ and $t \in \mathbb{R}$. Let $H_{1}, \ldots, H_{L}$ be Hermitian operators, and $\Lambda:=$ $\max _{j}\left\|H_{j}\right\|$. Then

$$
\begin{equation*}
\left\|\exp \left(-i t \sum_{j=1}^{L} H_{j}\right)-\left[\prod_{j=1}^{L} \exp \left(-\frac{i t}{r} H_{j}\right)\right]^{r}\right\| \leq \frac{(L \Lambda t)^{2}}{r} \exp \left(\frac{L \Lambda|t|}{r}\right) . \tag{61}
\end{equation*}
$$

Proof. Let $\lambda=-i t$. We start by considering $\left\|\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)-\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right\|$. Using Lemma F. 1 and Lemma F.2, we find

$$
\begin{align*}
\left\|\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)-\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right\| & =\left\|\mathcal{R}_{1}\left(\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)-\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right)\right\|  \tag{62}\\
& \leq\left\|\mathcal{R}_{1}\left(\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)\right)\right\|+\left\|\mathcal{R}_{1}\left(\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right)\right\|  \tag{63}\\
& \leq \mathcal{R}_{1}\left(\exp \left(|\lambda| \cdot\left|\sum_{j=1}^{L} H_{j}\right|\right)\right)+\mathcal{R}_{1}\left(\exp \left(\sum_{j=1}^{L}|\lambda| \cdot\left\|H_{j}\right\|\right)\right)  \tag{64}\\
& \leq 2 \mathcal{R}_{1}(\exp (|\lambda| L \Lambda))  \tag{65}\\
& \leq(|\lambda| L \Lambda)^{2} \exp (|\lambda| L \Lambda) . \tag{66}
\end{align*}
$$

We then divide the evolution into $r$ segments and apply the above inequality to each segment, giving

$$
\begin{align*}
\left\|\exp \left(-i t \sum_{j=1}^{L} H_{j}\right)-\left[\prod_{j=1}^{L} \exp \left(-\frac{i t}{r} H_{j}\right)\right]^{r}\right\| & \leq r\left\|\exp \left(-\frac{i t}{r} \sum_{j=1}^{L} H_{j}\right)-\prod_{j=1}^{L} \exp \left(-\frac{i t}{r} H_{j}\right)\right\|  \tag{67}\\
& \leq r\left(\frac{t L \Lambda}{r}\right)^{2} \exp \left(\frac{|t| L \Lambda}{r}\right)  \tag{68}\\
& =\frac{(L \Lambda t)^{2}}{r} \exp \left(\frac{L \Lambda|t|}{r}\right), \tag{69}
\end{align*}
$$

which completes the proof.
A similar bound holds for the higher-order case, as follows.
Proposition F.4. Let $r \in \mathbb{N}$ and $t \in \mathbb{R}$. Let $H_{1}, \ldots, H_{L}$ be Hermitian operators, and $\Lambda:=$ $\max _{j}\left\|H_{j}\right\|$. Then

$$
\begin{equation*}
\left\|\exp \left(-i t \sum_{j=1}^{L} H_{j}\right)-\left[S_{2 k}\left(-\frac{i t}{r}\right)\right]^{r}\right\| \leq \frac{\left(2 L 5^{k-1} \Lambda|t|\right)^{2 k+1}}{3 r^{2 k}} \exp \left(\frac{2 L 5^{k-1} \Lambda|t|}{r}\right) \tag{70}
\end{equation*}
$$

We omit the proof, which follows along the same lines as the proof of Proposition F.3.
When performing quantum simulation, our goal is to ensure that the error is at most some given $\epsilon$. Thus, to apply Proposition F. 3 and Proposition F.4, we must choose a value of $r$ that ensures the right-hand side is at most $\epsilon$. One approach [11] is to give a closed-form expression for some suitable $r$ as a function of $\epsilon$ (and other simulation parameters). We call the resulting error bound an analytic bound.

Definition 1. Let $t \in \mathbb{R}$ and $\epsilon>0$. Let $H_{1}, \ldots, H_{L}$ be Hermitian operators, and $\Lambda:=\max _{j}\left\|H_{j}\right\|$. Then the first-order analytic bound is given by

$$
\begin{equation*}
r_{1}=\left\lceil\max \left\{L|t| \Lambda, \frac{e(L t \Lambda)^{2}}{\epsilon}\right\}\right\rceil \tag{71}
\end{equation*}
$$

and the $(2 k)$ th order analytic bound is given by

$$
\begin{equation*}
r_{2 k}=\left\lceil\max \left\{2 L 5^{k-1} \Lambda|t|, \sqrt[2 k]{\frac{e\left(2 L 5^{k-1} \Lambda|t|\right)^{2 k+1}}{3 \epsilon}}\right\}\right\rceil \tag{72}
\end{equation*}
$$

Analytic bounds can be computed efficiently and result in circuits approximating the time evolution up to the required precision. However, because the error estimates are pessimistic, the corresponding circuits are unnecessarily large.

A slight improvement can be obtained by searching for the smallest possible $r$ that satisfies the constraints given by Proposition F. 3 and Proposition F.4. We call the resulting error bound a minimized bound.

Definition 2. Let $t \in \mathbb{R}$ and $\epsilon>0$. Let $H_{1}, \ldots, H_{L}$ be Hermitian operators, and $\Lambda:=\max _{j}\left\|H_{j}\right\|$. Then the first-order minimized bound is given by

$$
\begin{equation*}
r_{1}=\min \left\{r \in \mathbb{N}: \frac{(L \Lambda t)^{2}}{r} \exp \left(\frac{L \Lambda|t|}{r}\right) \leq \epsilon\right\} \tag{73}
\end{equation*}
$$

and the ( $2 k$ )th-order minimized bound is given by

$$
\begin{equation*}
r_{2 k}=\min \left\{r \in \mathbb{N}: \frac{\left(2 L 5^{k-1} \Lambda|t|\right)^{2 k+1}}{3 r^{2 k}} \exp \left(\frac{2 L 5^{k-1} \Lambda|t|}{r}\right) \leq \epsilon\right\} \tag{74}
\end{equation*}
$$

Minimized bounds can be computed efficiently using binary search. First, we compute the error bound with $r=1$. If the result is within our desired accuracy $\epsilon$, we stop; otherwise, we double $r$ and repeat the procedure. This allows us to bracket the value of $r$ in logarithmic time. Then we perform binary search to determine the precise value of $r$.

## F. 2 Commutator bounds

If all terms in the Hamiltonian commute, then even the lowest-order product formula has no error. Intuitively, if many pairs of terms commute, then the error should be small. Indeed, the secondorder contribution to the error in the first-order formula involves commutators between all pairs of terms, and this fact can be used to strengthen the bound [50].

In this section, we present improved error bounds that take advantage of commuting pairs of terms in the Hamiltonian. In particular, we give commutator bounds for the second- and fourthorder formulas, whose constructions can in principle be adapted to higher-order cases. To the best of our knowledge, these are the first error bounds for higher-order product formulas that take advantage of commutator structure.

In Section F.2.1, we present abstract versions of the commutator bounds, stated in terms of coefficients that count the number of terms in the Hamiltonian with particular commutation relations. We specialize these bounds to our particular system in Section F.2.2, using the symmetry of our Hamiltonian to substantially simplify application of the bounds.

## F.2.1 Abstract commutator bounds

We begin by presenting a bound for the first-order formula as described in [50], giving an explicit bound for the higher-order terms.

Theorem F. 5 (First-order commutator bound). Let $H_{1}, \ldots, H_{L}$ be Hermitian operators with norm at most $\Lambda:=\max _{j}\left\|H_{j}\right\|$. Let $C:=\left|\left\{\left(H_{i}, H_{j}\right):\left[H_{i}, H_{j}\right] \neq 0, i<j\right\}\right|$ be the number of noncommuting pairs of operators, and let $t \in \mathbb{R}$. Then

$$
\begin{equation*}
\left\|\exp \left(-i t \sum_{j=1}^{L} H_{j}\right)-\left[\prod_{j=1}^{L} \exp \left(-\frac{i t}{r} H_{j}\right)\right]^{r}\right\| \leq C \frac{(\Lambda t)^{2}}{r}+\frac{(L \Lambda|t|)^{3}}{3 r^{2}} \exp \left(\frac{L \Lambda|t|}{r}\right) \tag{75}
\end{equation*}
$$

Proof. We show that

$$
\begin{equation*}
\left\|\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)-\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right\| \leq C(|\lambda| \Lambda)^{2}+\frac{(|\lambda| L \Lambda)^{3}}{3} \exp (|\lambda| L \Lambda), \tag{76}
\end{equation*}
$$

which implies the claimed result by the triangle inequality. The upper bound (76) can be established by explicitly computing the second-order error and bounding the higher-order errors by the norm of the remainder $\mathcal{R}_{2}$. The second-order error is bounded as follows:

$$
\begin{align*}
& \left\|\frac{\lambda^{2}}{2}\left(\sum_{j=1}^{L} H_{j}\right)^{2}-\left[\frac{\lambda^{2}}{2} \sum_{j=1}^{L} H_{j}^{2}+\lambda^{2} \sum_{j<k} H_{j} H_{k}\right]\right\| \\
& \quad=\left\|\left[\frac{\lambda^{2}}{2} \sum_{j=1}^{L} H_{j}^{2}+\frac{\lambda^{2}}{2} \sum_{j<k} H_{j} H_{k}+\frac{\lambda^{2}}{2} \sum_{j>k} H_{j} H_{k}\right]-\left[\frac{\lambda^{2}}{2} \sum_{j=1}^{L} H_{j}^{2}+\lambda^{2} \sum_{j<k} H_{j} H_{k}\right]\right\|  \tag{77}\\
& \quad=\left\|\frac{\lambda^{2}}{2}\left[\sum_{j>k} H_{j} H_{k}-\sum_{j<k} H_{j} H_{k}\right]\right\|=\left\|\frac{\lambda^{2}}{2}\left[\sum_{j>k} H_{j} H_{k}-H_{k} H_{j}\right]\right\|  \tag{78}\\
& \quad \leq C|\lambda|^{2} \Lambda^{2} . \tag{79}
\end{align*}
$$

The rest of the proof proceeds similarly to the second half of the proof of Proposition F.3; we omit the details.

Theorem F. 6 (Second-order commutator bound). Let $H_{1}, \ldots, H_{L}$ be Hermitian operators with norm at most $\Lambda:=\max _{j}\left\|H_{j}\right\|$, where each $H_{j}$ is a tensor product of Pauli operators. Define the augmented set of Hamiltonians

$$
\tilde{H}_{j}= \begin{cases}H_{j}, & 1 \leq j \leq L  \tag{80}\\ H_{2 L-j+1}, & L+1 \leq j \leq 2 L\end{cases}
$$

Let $f(i, j)=1$ if $\tilde{H}_{i}, \tilde{H}_{j}$ commute and $f(i, j)=-1$ otherwise. Finally, let

$$
\begin{align*}
D: & =|\{(i, j): f(i, j)=-1, i \neq j\}|,  \tag{81}\\
T_{1}:= & |\{(i, j, k): f(i, j)=f(j, k)=f(i, k)=1, i<j<k\}|,  \tag{82}\\
T_{2}: & =|\{(i, j, k): f(i, j)=1, f(j, k)=f(i, k)=-1, i<j<k\}| \\
& +|\{(i, j, k): f(i, j)=f(i, k)=-1, f(j, k)=1, i<j<k\}|,  \tag{83}\\
T_{3}: & =|\{(i, j, k): f(i, j)=f(j, k)=-1, f(i, k)=1, i<j<k\}|,  \tag{84}\\
T_{4}: & =\mid\{(i, j, k): \text { all other cases }\} \mid \tag{85}
\end{align*}
$$

where $i, j, k \in\{1, \ldots, 2 L\}$, and let $t \in \mathbb{R}$. Then

$$
\begin{align*}
& \left\|\exp \left(-i t \sum_{j=1}^{L} H_{j}\right)-\left[S_{2}\left(-\frac{i t}{r}\right)\right]^{r}\right\| \\
& \quad \leq \frac{\Lambda^{3}|t|^{3}}{r^{2}}\left\{\frac{1}{24} D+\frac{1}{12} T_{2}+\frac{1}{6} T_{3}+\frac{1}{8} T_{4}\right\}+\frac{4(L \Lambda t)^{4}}{3 r^{3}} \exp \left(\frac{2 L \Lambda|t|}{r}\right) \tag{86}
\end{align*}
$$

Proof. As in the proof of Theorem F.5, we explicitly compute the third-order error and bound the higher-order terms by $\mathcal{R}_{3}$. First we show that the first-order formula

$$
\begin{equation*}
\left\|\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)-\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)\right\| \tag{87}
\end{equation*}
$$

has a third-order error of at most

$$
\begin{equation*}
|\lambda|^{3} \Lambda^{3}\left(\frac{1}{3} \bar{D}+\frac{2}{3} \bar{T}_{2}+\frac{4}{3} \bar{T}_{3}+\bar{T}_{4}\right), \tag{88}
\end{equation*}
$$

where the coefficients $\bar{D}, \bar{T}_{2}, \bar{T}_{3}, \bar{T}_{4}$ are defined as in (81)-(85), but with respect to the original Hamiltonians $\left\{H_{j}\right\}_{j=1}^{L}$ instead of $\left\{\tilde{H}_{j}\right\}_{j=1}^{2 L}$.

The third-order term in $\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)$ is

$$
\begin{equation*}
\frac{1}{3!}\left(\lambda \sum_{j=1}^{L} H_{j}\right)^{3}=\frac{\lambda^{3}}{6} \sum_{i, j, k} H_{i} H_{j} H_{k} \tag{89}
\end{equation*}
$$

whereas the third-order term in $\prod_{j=1}^{L} \exp \left(\lambda H_{j}\right)$ is

$$
\begin{align*}
& \frac{\lambda^{3}}{6} \sum_{i} H_{i}^{3}+\frac{\lambda^{3}}{2} \sum_{i<k} H_{i}^{2} H_{k}+\frac{\lambda^{3}}{2} \sum_{i<k} H_{i} H_{k}^{2}+\lambda^{3} \sum_{i<j<k} H_{i} H_{j} H_{k} \\
& \quad=\frac{\lambda^{3}}{6} \sum_{i} H_{i}^{3}+\frac{\lambda^{3}}{2} \sum_{i \neq k} H_{i}^{2} H_{k}+\lambda^{3} \sum_{i<j<k} H_{i} H_{j} H_{k} \tag{90}
\end{align*}
$$

where we have used the fact that the square of any Pauli operator is the identity. Taking the difference gives

$$
\begin{equation*}
\frac{\lambda^{3}}{6} \sum_{i \neq j} H_{i}\left[H_{j}, H_{i}\right]+\frac{\lambda^{3}}{6} \sum_{\substack{i, j, k \\ \text { pairwise different }}} H_{i} H_{j} H_{k}-\lambda^{3} \sum_{i<j<k} H_{i} H_{j} H_{k} \tag{91}
\end{equation*}
$$

The norm of the first term is at most

$$
\begin{equation*}
\frac{1}{3}|\lambda|^{3} \Lambda^{3} \bar{D} \tag{92}
\end{equation*}
$$

whereas the last two terms can be written as follows:

$$
\begin{align*}
& \frac{\lambda^{3}}{6} \sum_{\substack{i_{1}, i_{2}, i_{3} \\
\text { pairwis differt }}} H_{i_{1}} H_{i_{2}} H_{i_{3}}-\lambda^{3} \sum_{i_{1}<i_{2}<i_{3}} H_{i_{1}} H_{i_{2}} H_{i_{3}} \\
& =\frac{\lambda^{3}}{6} \sum_{i_{1}<i_{2}<i_{3}} \sum_{\sigma \in \operatorname{Sym}(3)} H_{i_{\sigma(1)}} H_{i_{\sigma(2)}} H_{i_{\sigma(3)}}-\lambda^{3} \sum_{i_{1}<i_{2}<i_{3}} H_{i_{1}} H_{i_{2}} H_{i_{3}}  \tag{93}\\
& =\frac{\lambda^{3}}{6} \sum_{i_{1}<i_{2}<i_{3}}(1+f(1,2)+f(2,3)+f(1,2) f(1,3) f(2,3)  \tag{94}\\
& \quad+f(1,3) f(1,2)+f(1,3) f(2,3)) H_{i_{1}} H_{i_{2}} H_{i_{3}}-\lambda^{3} \sum_{i_{1}<i_{2}<i_{3}} H_{i_{1}} H_{i_{2}} H_{i_{3}} . \tag{95}
\end{align*}
$$

(Here $\operatorname{Sym}(3)$ denotes the symmetric group on three elements.) By performing case analysis, we can evaluate the coefficients and upper bound the norm by

$$
\begin{equation*}
|\lambda|^{3} \Lambda^{3}\left(\frac{2}{3} \bar{T}_{2}+\frac{4}{3} \bar{T}_{3}+\bar{T}_{4}\right) \tag{96}
\end{equation*}
$$

Combining (92) and (96), we obtain the claimed upper bound (88) for the third-order error in the first-order formula.

Now we consider the second-order formula. Similarly to the proof of Theorem F.5, we begin by proving the bound

$$
\begin{equation*}
\left\|\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)-S_{2}(\lambda)\right\| \leq(|\lambda| \Lambda)^{3}\left(\frac{D}{24}+\frac{T_{2}}{12}+\frac{T_{3}}{6}+\frac{T_{4}}{8}\right)+\frac{4}{3}(L|\lambda| \Lambda)^{4} \exp (2 L|\lambda| \Lambda), \tag{97}
\end{equation*}
$$

which implies (86) by the triangle inequality.
To establish (97), we apply (88) to the augmented Hamiltonian list $\left\{\tilde{H}_{j}\right\}_{j=1}^{2 L}$ with $\lambda$ replaced by $\frac{\lambda}{2}$. This shows that the third-order error is at most

$$
\begin{equation*}
\frac{|\lambda|^{3} \Lambda^{3}}{8}\left(\frac{D}{3}+\frac{2 T_{2}}{3}+\frac{4 T_{3}}{3}+T_{4}\right) . \tag{98}
\end{equation*}
$$

The higher-order errors can be bounded by a routine calculation as

$$
\begin{equation*}
\mathcal{R}_{3}\left(\exp \left(\lambda \sum_{j=1}^{L} H_{j}\right)-S_{2}(\lambda)\right) \leq 2 \mathcal{R}_{3}(\exp (2 L|\lambda| \Lambda)) \leq 2 \frac{(2 L|\lambda| \Lambda)^{4}}{4!} \exp (2 L|\lambda| \Lambda) . \tag{99}
\end{equation*}
$$

This completes the proof of (97). The remainder of the proof proceeds similarly to the second half of the proof of Proposition F.3.

A similar bound holds for the fourth-order formula, as follows.
Theorem F. 7 (Fourth-order commutator bound). Let $H_{1}, \ldots, H_{L}$ be Hermitian operators with norm at most $\Lambda:=\max _{j}\left\|H_{j}\right\|$, where each $H_{j}$ is a tensor product of Pauli operators. Define the augmented set of Hamiltonians

$$
\tilde{H}_{j}=\left\{\begin{array}{ll}
H_{j-2 h L}, & 2 h L+1 \leq j \leq(2 h+1) L  \tag{100}\\
H_{2(h+1) L-j+1}, & (2 h+1) L+1 \leq j \leq 2(h+1) L
\end{array} \quad h \in\{0,1,2,3,4\} .\right.
$$

Let $f(i, j)=1$ if $\tilde{H}_{i}, \tilde{H}_{j}$ commute and $f(i, j)=-1$ otherwise. Finally, let

$$
\begin{align*}
N_{a}:= & \left|\left\{(i, j): f(i, j)=a_{1}, i<j\right\}\right|  \tag{101}\\
N_{b_{1} b_{2} b_{3}}:= & \left|\left\{(i, j, k): f(i, j)=b_{1}, f(i, k)=b_{2}, f(j, k)=b_{3}, i<j<k\right\}\right|  \tag{102}\\
N_{c_{1} \ldots c 6}:= & \mid\left\{(i, j, k, l): f(i, j)=c_{1}, f(i, k)=c_{2}, f(i, l)=c_{3},\right. \\
& \left.f(j, k)=c_{4}, f(j, l)=c_{5}, f(k, l)=c_{6}, i<j<k<l\right\} \mid  \tag{103}\\
N_{d_{1} \ldots d_{10}}:= & \mid\left\{(i, j, k, l, m): f(i, j)=d_{1}, f(i, k)=d_{2}, f(i, l)=d_{3},\right. \\
& f(i, m)=d_{4}, f(j, k)=d_{5}, f(j, l)=d_{6}, f(j, m)=d_{7}, f(k, l)=d_{8}, \\
& \left.f(k, m)=d_{9}, f(l, m)=d_{10}, i<j<k<l<m\right\} \mid \tag{104}
\end{align*}
$$

where $i, j, k, l, m \in\{1, \ldots, 10 L\}$, let $t \in \mathbb{R}$, and let $p:=1 /\left(4-4^{1 / 3}\right)$. Then

$$
\begin{align*}
\left\|\exp \left(-i t \sum_{j=1}^{L} H_{j}\right)-\left[S_{4}\left(-\frac{i t}{r}\right)\right]^{r}\right\| \leq & \left(\frac{4 p-1}{2} \Lambda|t|\right)^{5} \frac{1}{5!r^{4}}\left\{\sum_{a \in \pm 1} C_{a} N_{a}+\sum_{b_{1}, b_{2}, b_{3} \in \pm 1} C_{b_{1} b_{2} b_{3}} N_{b_{1} b_{2} b_{3}}\right. \\
& \left.+\sum_{c_{1}, \ldots, c_{6} \in \pm 1} C_{c_{1} \ldots c_{6}} N_{c_{1} \ldots c_{6}}+\sum_{d_{1}, \ldots, d_{10} \in \pm 1} C_{d_{1} \ldots d_{10}} N_{d_{1} \ldots d_{10}}\right\} \\
& +2 \frac{(5(4 p-1) L \Lambda t)^{6}}{6!\cdot r^{5}} \exp \left(\frac{(5(4 p-1) L \Lambda|t|}{r}\right) \tag{105}
\end{align*}
$$

for some real coefficients $C_{a}, C_{b_{1} b_{2} b_{3}}, C_{c_{1} \ldots c_{6}}, C_{d_{1} \ldots d_{10}}$.
We omit the proof, which proceeds along similar lines to that of Theorem F.6. Note that similar bounds also hold for higher-order formulas.

The coefficients $C_{a}, C_{b_{1} b_{2} b_{3}}, C_{c_{1} \ldots c_{6}}, C_{d_{1} \ldots d_{10}}$ can in principle be determined by a computer program. The list of coefficients is long, so we omit it here. In Section F.2.2, we give a specialized version of this theorem that can be applied to the Hamiltonian (1) without explicitly computing all the coefficients.

To give an idea of how the bound is proved, we show how to determine the coefficient $C_{-1}$ in (105). Similar arguments can be used to determine all the coefficients in the bound.

First consider the fifth-order terms of the expression

$$
\begin{equation*}
\exp \left(\sum_{j=1}^{L} H_{j} \lambda\right)-\exp \left(H_{1} \lambda\right) \cdots \exp \left(H_{L} \lambda\right) . \tag{106}
\end{equation*}
$$

The coefficient $C_{-1}$ of $N_{-1}$ counts the pairs of non-commuting terms $H_{i}$ and $H_{j}$. The second term in (106) contributes

$$
\begin{equation*}
\frac{\lambda^{4}}{4!} \lambda \sum_{i<j}\left(H_{i}^{4} H_{j}+H_{i} H_{j}^{4}\right)+\frac{\lambda^{3}}{3!} \frac{\lambda^{2}}{2!} \sum_{i<j}\left(H_{i}^{3} H_{j}^{2}+H_{i}^{2} H_{j}^{3}\right), \tag{107}
\end{equation*}
$$

whereas the first term in (106) contributes

$$
\begin{align*}
& \quad \frac{\lambda^{5}}{5!} \sum_{i \neq j}\left(H_{i}^{4} H_{j}+H_{i}^{3} H_{j} H_{i}+H_{i}^{2} H_{j} H_{i}^{2}+H_{i} H_{j} H_{i}^{3}+H_{j} H_{i}^{4}\right) \\
& +\frac{\lambda^{5}}{5!} \sum_{i \neq j}\left(H_{i}^{3} H_{j}^{2}+H_{i}^{2} H_{j} H_{i} H_{j}+H_{i}^{2} H_{j}^{2} H_{i}+H_{i} H_{j} H_{i}^{2} H_{j}+H_{i} H_{j} H_{i} H_{j} H_{i}\right.  \tag{108}\\
& \left.\quad+H_{i} H_{j}^{2} H_{i}^{2}+H_{j} H_{i}^{3} H_{j}+H_{j} H_{i}^{2} H_{j} H_{i}+H_{j} H_{i} H_{j} H_{i}^{2}+H_{j}^{2} H_{i}^{3}\right) .
\end{align*}
$$

Since we assume the terms of the Hamiltonian are tensor products of Pauli operators, we can interchange the order of multiplication, possibly introducing minus signs. Thus (108) equals

$$
\begin{align*}
& \frac{\lambda^{5}}{5!} \sum_{i<j}\left((3+2 f(i, j)) H_{i}^{4} H_{j}+(3+2 f(i, j)) H_{i} H_{j}^{4}\right) \\
& \quad+\frac{\lambda^{5}}{5!} \sum_{i<j}\left((6+4 f(i, j)) H_{i}^{3} H_{j}^{2}+(6+4 f(i, j)) H_{i}^{2} H_{j}^{3}\right) . \tag{109}
\end{align*}
$$

Subtracting (109) from (107) and setting $f(i, j)=-1$, we find

$$
\begin{equation*}
\frac{\lambda^{5}}{5!} \sum_{i<j}\left(4 H_{i}^{4} H_{j}+4 H_{i} H_{j}^{4}+8 H_{i}^{3} H_{j}^{2}+8 H_{i}^{2} H_{j}^{3}\right) \tag{110}
\end{equation*}
$$

whose spectral norm is bounded by

$$
\begin{equation*}
\frac{|\lambda|^{5}}{5!} 24 \Lambda^{5} N_{-1} . \tag{111}
\end{equation*}
$$

Comparing the result to (105), we find that $C_{-1}=24$.

## F.2.2 Concrete commutator bounds

To apply the above commutator bounds, we must compute the number of tuples of terms in the Hamiltonian satisfying certain commutation relations (e.g., equations (101)-(104) for the fourthorder bound). While this can be done in polynomial time provided the degree is constant, a direct approach is prohibitive in practice. For example, consider the problem of counting 5-tuples of $\Theta(n)$ terms, as required for the fourth-order commutator bound. A naive iteration over all 5 -tuples takes time $O\left(n^{5}\right)$, and the multiplicative constant is large since our Hamiltonian gives rise to $40 n$ terms $\tilde{H}_{i}$. Thus a direct computation can be infeasible even for small $n$.

However, for the Hamiltonian (1), we show that each number of tuples is given by a lowdegree polynomial in $n$. In turn, this means that the lowest-order contribution to the error is also a polynomial in $n$. Thus, by performing polynomial interpolation on a constant number of numerically-obtained values, we can determine a closed-form expression for general $n$. A similar result holds for other Hamiltonians with appropriate symmetries, so this approach could be applied more broadly.

We first assume the Hamiltonian acts locally on a line (in particular, we assume open boundary conditions instead of periodic ones). For example, we could consider the Hamiltonian

$$
\begin{equation*}
\sum_{j=1}^{n-1}\left(\vec{\sigma}_{j} \cdot \vec{\sigma}_{j+1}+h_{j} \sigma_{j}^{z}\right) \tag{112}
\end{equation*}
$$

which is simply the Heisenberg model (1) without the boundary term, and without the local magnetic field on the rightmost spin (that term could also be handled, but we exclude it to simplify the presentation). We show that if a counting problem for such a Hamiltonian is oblivious (defined below), then the result is a polynomial in $n$. We then show that a similar property holds for Hamiltonians with periodic boundary conditions, giving a method for evaluating the commutator bound for our model system.

We focus on the list of terms in the Hamiltonian to be applied in each segment of a product formula. Specifically, consider $n$ qubits arranged on a line, and suppose we are given Hermitian operators $G_{1}, \ldots, G_{\nu}$ with values $d_{1}, \ldots, d_{\nu} \in\{-1,1\}$ indicating shift directions. We assume that
each term $G_{i}$ acts nontrivially only on the first $k$ qubits if $d_{i}=1$ or on the last $k$ qubits if $d_{i}=-1$. This determines an ordered list of length $\mu:=\nu(n-k+1)$ including all translations of the given terms, of the form

$$
\begin{align*}
& \left(G_{1}, R^{d_{1}} G_{1} R^{-d_{1}}, \ldots, R^{d_{1}(n-k)} G_{1} R^{-d_{1}(n-k)}\right. \\
& G_{2}, R^{d_{2}} G_{2} R^{-d_{2}}, \ldots, R^{d_{2}(n-k)} G_{2} R^{-d_{2}(n-k)} \\
& \vdots \\
& \left.G_{\nu}, R^{d_{\nu}} G_{\nu} R^{-d_{\nu}}, \ldots, R^{d_{\nu}(n-k)} G_{\nu} R^{-d_{\nu}(n-k)}\right) \tag{113}
\end{align*}
$$

where $R$ is the operator that shifts the qubits right by one position. For each $i \in \llbracket \mu \rrbracket:=\{1, \ldots, \mu\}$, let $H_{i}$ denote the $i$ th term in this list.

In the example (112), the $n$-qubit Hamiltonian consists of $4(n-1)$ terms. We adopt a specific ordering of the terms, namely

$$
\begin{equation*}
\sigma_{1}^{x} \sigma_{2}^{x}, \ldots, \sigma_{n-1}^{x} \sigma_{n}^{x}, \sigma_{1}^{y} \sigma_{2}^{y}, \ldots, \sigma_{n-1}^{y} \sigma_{n}^{y}, \sigma_{1}^{z} \sigma_{2}^{z}, \ldots, \sigma_{n-1}^{z} \sigma_{n}^{z}, \sigma_{1}^{z}, \ldots, \sigma_{n-1}^{z} \tag{114}
\end{equation*}
$$

Using the fourth-order product formula, every term of the Hamiltonian repeats 10 times in a fixed order within each segment. Thus, we obtain an ordered list of 2-local operators (i.e., $k=2$ ) with length $\mu=40(n-1)$, and we see that it has the form of (113) with $\nu=40$.

We now introduce two mappings, supp and ival. The mapping supp: $\llbracket \mu \rrbracket \rightarrow \mathbb{Z}_{2}^{n}$ returns the support of a term in the Hamiltonian. The support of $G_{i}$ is 1 in the first $k$ positions and 0 in the rest if $d_{i}=1$, or 1 in the last $k$ positions and 0 in the rest if $d_{i}=-1$. Similarly, the support of $R^{j} G_{i} R^{-j}$ has a 1 at positions $j+1, j+2, \ldots, j+k$ if $d_{i}=1$, or a 1 at positions $n-j-k+1, n-j-k+2, \ldots, n-j$ if $d_{i}=-1$. Note that the support may include positions at which the Hamiltonian acts trivially.

To study the action of a $\tau$-tuple of Hamiltonian operators, we take the entrywise or of all elements of the tuple, defining

$$
\begin{align*}
\text { supp: } \mathcal{S}_{\tau}^{\mu} & \rightarrow \mathbb{Z}_{2}^{n}  \tag{115}\\
\operatorname{supp}\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right) & =\bigvee_{j=1}^{\tau} \operatorname{supp}\left(i_{j}\right) \tag{116}
\end{align*}
$$

Here $\mathcal{S}_{\tau}^{\mu}:=\{S \subseteq \llbracket \mu \rrbracket:|S|=\tau\}$ is the set of all subsets of $\llbracket \mu \rrbracket$ with size $\tau$. Since the support of each term in the Hamiltonian has Hamming weight $k$, the Hamming weight of $\operatorname{supp}\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)$ is at most $\tau k$.

The mapping ival: $\mathbb{Z}_{2}^{n} \rightarrow \mathbb{Z}_{\tau k}^{\tau}$ returns the interval pattern for a given subset of qubits. The interval pattern is a list of the numbers of consecutive 1 s in the string, padded with trailing 0 s if necessary to have length $\tau k$. For example, ival $\left(\left[\begin{array}{llll}1 & 0 & 1 & 1\end{array} 0\right]\right)=\left[\begin{array}{lllll}1 & 2 & 0 & \ldots & 0\end{array}\right]$.

Recall that, by assumption, every term in the Hamiltonian acts nontrivially on at most $k$ consecutive qubits, and its support is defined to include all $k$ relevant positions (possibly including ones where the Hamiltonian acts trivially). Thus any such term can occupy at most one interval of any given interval pattern.

Our main goal is to count the number of $\tau$-tuples of Hamiltonian operators satisfying a given commutation pattern. We regard this as a special case of a more general task, which we call oblivious counting.

Let $f: \mathcal{S}_{\tau}^{\mu} \rightarrow \mathbb{Z}_{2}$ be a binary-valued function that represents whether a tuple of operators has some specified property. For example, when computing the fourth-order commutator bound, we
set $f\left(\left\{i_{1}, i_{2}, i_{3}, i_{4}, i_{5}\right\}\right)=1$ if the tuple $\left(H_{i_{1}}, \ldots, H_{i_{5}}\right)$ satisfies a given commutation pattern, and $f\left(\left\{i_{1}, i_{2}, i_{3}, i_{4}, i_{5}\right\}\right)=0$ otherwise. Our task now is to count the number of tuples satisfying the property, i.e., to compute

$$
\begin{equation*}
\left|f^{-1}(1)\right|=\left|\left\{\left\{i_{1}, \ldots, i_{\tau}\right\}: f\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=1\right\}\right| . \tag{117}
\end{equation*}
$$

We say the counting is oblivious if, for any fixed $z \in \mathbb{Z}_{\tau k}^{\tau}$, the quantity

$$
\begin{equation*}
\left|\left\{\left\{i_{1}, \ldots, i_{\tau}\right\}: f\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=1, \operatorname{supp}\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=y, \operatorname{ival}(y)=z\right\}\right| \tag{118}
\end{equation*}
$$

is identical for all $n$ and all $y \in \mathbb{Z}_{2}^{n}$ with ival $(y)=z$. In other words, an oblivious counting problem depends only on the interval pattern of the $\tau$-tuple, not on its precise support or the size of the underlying system.

Our main result for counting such tuples is as follows.
Theorem F.8. Let $G_{1}, \ldots, G_{\nu}$ be Hermitian operators with $d_{1}, \ldots, d_{\nu} \in\{-1,1\}$. We assume that each term $G_{i}$ acts as the identity on all but the first (if $d_{i}=1$ ) or last (if $d_{i}=-1$ ) $k$ qubits. Suppose that $f: \mathcal{S}_{\tau}^{\mu} \rightarrow \mathbb{Z}_{2}$ induces an oblivious counting problem, as defined above. Then $\left|f^{-1}(1)\right|$ is a polynomial in $n$.

Proof. Our task is to compute

$$
\begin{align*}
\left|f^{-1}(1)\right| & =\left|\left\{\left\{i_{1}, \ldots, i_{\tau}\right\}: f\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=1\right\}\right|  \tag{119}\\
& =\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} \sum_{y \in \mathbb{Z}_{2}^{n}} C_{z y}(n), \tag{120}
\end{align*}
$$

where

$$
\begin{equation*}
C_{z y}(n):=\left|\left\{\left\{i_{1}, \ldots, i_{\tau}\right\}: f\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=1, \operatorname{supp}\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=y, \operatorname{ival}(y)=z\right\}\right| . \tag{121}
\end{equation*}
$$

We claim that
(i) a constant number of values $z \in \mathbb{Z}_{\tau k}^{\tau}$ (independent of $n$ ) contribute to the sum; and
(ii) for fixed $z \in \mathbb{Z}_{\tau k}^{\tau}$, the number of $y \in \mathbb{Z}_{2}^{n}$ such that $\operatorname{ival}(y)=z$ is a polynomial in $n$.

Claim (i) follows because we assume the size $\tau$ of each tuple and the number $k$ of local operators are both independent of $n$. To prove claim (ii), fix an interval pattern

$$
z=\left[\begin{array}{llllll}
z_{1} & \ldots & z_{r} & 0 & \ldots & 0 \tag{122}
\end{array}\right] \in \mathbb{Z}_{\tau k}^{\tau}
$$

and define

$$
\begin{equation*}
T\left(n ; z_{1}, \ldots, z_{r}\right):=\left|\left\{y \in \mathbb{Z}_{2}^{n}: \operatorname{ival}(y)=z\right\}\right| . \tag{123}
\end{equation*}
$$

We show by induction on $r$ that $T\left(n ; z_{1}, \ldots, z_{r}\right)$ is a polynomial of degree at most $r$. The base case requires us to place an interval of length $z_{1}$ in $n$ slots, and clearly $T\left(n ; z_{1}\right)=n-z_{1}+1$, a polynomial of degree 1 .

For the induction step, suppose the claim is true for any $r-1$ intervals $\left(z_{1}, \ldots, z_{r-1}, 0, \ldots, 0\right)$. Using the hypothesis for the case of $r-1$ intervals, we find

$$
\begin{align*}
T\left(n ; z_{1}, \ldots, z_{r}\right)= & \sum_{j=1}^{n-z_{r}} T\left(n-z_{r}-j ; z_{1}, \ldots, z_{r-1}\right)  \tag{124}\\
= & \sum_{j=1}^{n-z_{r}} c_{r-1}\left(z_{1}, \ldots, z_{r-1}\right)\left(n-z_{r}-j\right)^{r-1} \\
& \quad+\cdots+c_{0}\left(z_{1}, \ldots, z_{r-1}\right)\left(n-z_{r}-j\right)^{0}, \tag{125}
\end{align*}
$$

where in the second equality we use the hypothesis to conclude that

$$
\begin{equation*}
T\left(n ; z_{1}, \ldots, z_{r-1}\right)=c_{r-1}\left(z_{1}, \ldots, z_{r-1}\right) n^{r-1}+\cdots+c_{0}\left(z_{1}, \ldots, z_{r-1}\right) n^{0} \tag{126}
\end{equation*}
$$

is a polynomial of degree at most $r-1$. It follows from Faulhaber's formula that $T\left(n ; z_{1}, \ldots, z_{r}\right)$ is a polynomial of degree at most $r$, which completes the inductive step.

To complete the proof of the main result, we observe that

$$
\begin{align*}
\left|f^{-1}(1)\right| & =\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} \sum_{\substack{y \in \mathbb{Z}_{2}^{n} \\
\text { ival }(y)=z}} C_{z y}(n)  \tag{127}\\
& =\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} p_{z}(n) C_{z}, \tag{128}
\end{align*}
$$

where, for each $z \in \mathbb{Z}_{\tau k}^{\tau}, p_{z}(n)$ is a polynomial in $n$ (by claim (ii)). Here we denote $C_{z}:=C_{z y}(n)$ since it is independent of $n$ and $y$, because we assume the counting problem is oblivious.

The above theorem relies crucially on the assumption that the terms in the Hamiltonian are arranged on a line. We now describe a variant of the theorem that handles Hamiltonians with periodic boundary conditions. To this end, consider $n$ qubits arranged on a ring, and again suppose we are given Hermitian operators $G_{1}, \ldots, G_{\nu}$ with values $d_{1}, \ldots, d_{\nu} \in\{-1,1\}$ indicating shift directions. We assume that each term $G_{i}$ acts as the identity except on the first $k$ qubits if $d_{i}=1$, or on the first $k-1$ qubits and the last qubit if $d_{i}=-1$. This determines an ordered list of length $\mu:=\nu n$ including all translations of the given terms, of the form

$$
\begin{align*}
& \left(G_{1}, R^{d_{1}} G_{1} R^{-d_{1}}, \ldots, R^{d_{1}(n-1)} G_{1} R^{-d_{1}(n-1)}\right. \\
& G_{2}, R^{d_{2}} G_{2} R^{-d_{2}}, \ldots, R^{d_{2}(n-1)} G_{2} R^{-d_{2}(n-1)} \\
& \vdots \\
& \left.G_{\nu}, R^{d_{\nu}} G_{\nu} R^{-d_{\nu}}, \ldots, R^{d_{\nu}(n-1)} G_{\nu} R^{-d_{\nu}(n-1)}\right) \tag{129}
\end{align*}
$$

where $R$ is the operator that shifts the qubits right by one position. For each $i \in \llbracket \mu \rrbracket:=\{1, \ldots, \mu\}$, let $H_{i}$ denote the $i$ th term in this list.

We define the support similarly as before, except that we take into account the periodic boundary conditions. In particular, the support of $R^{j} G_{i} R^{-j}$ with $d_{i}=1$ has a 1 at positions $(j \bmod n)+1,((j+1) \bmod n)+1, \ldots,((j+k-1) \bmod n)+1$. Likewise, the support of $R^{-j} G_{i} R^{j}$ with $d_{i}=-1$ has a 1 at positions $((n-j-1) \bmod n)+1,((n-j) \bmod n)+1, \ldots,((n-j+k-2) \bmod n)+1$. Again we extend the definition of support to a $\tau$-tuple of terms by taking the entrywise OR of all elements of the tuple. The notion of an interval pattern remains unchanged.

Our analysis divides the problem into a part where the boundary is irrelevant (so that we can apply the previous analysis) and a part where the boundary plays a role (which only involves a constant number of terms). To do this, we divide a support into three parts: the prefix, the suffix, and the infix.

Let $y \in \mathbb{Z}_{2}^{n}$ be the support of a tuple of operators. If $y_{1}=0$, we define prefix $(y) \in \mathbb{Z}_{2}^{n}$ to be the zero vector. Otherwise, $\operatorname{prefix}(y)$ is the indicator function of the first interval in $y$ (i.e., it is a $0 / 1$ vector that is 1 only on the first interval in $y$ ). Likewise, if $y_{n}=0$, we define suffix $(y) \in \mathbb{Z}_{2}^{n}$ to be the zero vector. Otherwise, suffix $(y)$ is the indicator function of the last interval in $y$. Then $\operatorname{infix}(y) \in \mathbb{Z}_{2}^{n}$ is the vector obtained by removing the first and the last interval from $y$, i.e.,

$$
\begin{equation*}
\operatorname{infix}(y):=y \oplus \operatorname{prefix}(y) \oplus \operatorname{suffix}(y) . \tag{130}
\end{equation*}
$$

Let $z \in \mathbb{Z}_{\tau k}^{\tau}$ be an interval pattern. We define the boundary as a mapping bndry: $\mathbb{Z}_{\tau k}^{\tau} \rightarrow \mathbb{Z}_{2}^{n}$. If there exists a $y \in \mathbb{Z}_{2}^{n}$ with $y_{1}=y_{n}=1$ such that $\operatorname{ival}(y)=z$, then $\operatorname{bndry}(z)=\operatorname{prefix}(y) \oplus \operatorname{suffix}(y)$. Otherwise, we set $\operatorname{bndry}(z)$ to be the zero vector. The mapping bndry is a well-defined function of the interval pattern $z$ since it does not depend on the specific choice of support $y$ (provided $y_{1}=y_{n}=1$, which is the only case in which we will use this function).

We now consider the counting problem. As before, we let $f: \mathcal{S}_{\tau}^{\mu} \rightarrow \mathbb{Z}_{2}$ be a binary-valued function that represents whether a tuple of operators has some specified property, and our goal is to compute $\left|f^{-1}(1)\right|$ as in (117). For a Hamiltonian on a ring with periodic boundary conditions, we say that the counting problem is oblivious if for any fixed $z \in \mathbb{Z}_{\tau k}^{\tau}$, the following two requirements are satisfied:
(O1) the quantity

$$
\begin{equation*}
\left|\left\{\left\{i_{1}, \ldots, i_{\tau}\right\}: f\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=1, \operatorname{supp}\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=y, \operatorname{ival}(y)=z\right\}\right| \tag{131}
\end{equation*}
$$

is identical for all $n$ and all $y \in \mathbb{Z}_{2}^{n}$ with $\operatorname{ival}(y)=z$ and $y_{1} y_{n}=0$; and
(O2) fixing $\omega \in \llbracket \tau \rrbracket$ and $\left\{j_{1}, \ldots, j_{\omega}\right\} \in \mathcal{S}_{\omega}^{\mu}$ with $\operatorname{supp}\left(\left\{j_{1}, \ldots, j_{\omega}\right\}\right)=\operatorname{bndry}(z)$, the quantity

$$
\begin{equation*}
\left|\left\{\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}: f_{j_{1}, \ldots, j_{\omega}}\left(\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}\right)=1, \operatorname{supp}\left(\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}\right)=\operatorname{infix}(y), \operatorname{ival}(y)=z\right\}\right| \tag{132}
\end{equation*}
$$

is identical for all $n$ and all $y \in \mathbb{Z}_{2}^{n}$ with $\operatorname{ival}(y)=z$ and $y_{1}=y_{n}=1$. Here $f_{j_{1}, \ldots, j_{\omega}}$ is a restricted version of $f$ defined by

$$
\begin{equation*}
f_{j_{1}, \ldots, j_{\omega}}\left(\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}\right)=f\left(\left\{j_{1}, \ldots, j_{\omega}, k_{1}, \ldots, k_{\tau-\omega}\right\}\right) . \tag{133}
\end{equation*}
$$

Note that the set of indices $\left\{j_{1}, \ldots, j_{\omega}\right\}$ is disjoint from $\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}$ as $\operatorname{bndry}(z)$ and $\operatorname{infix}(y)$ do not overlap. Therefore, the function $f_{j_{1}, \ldots, j_{\omega}}$ is well-defined. Note also that in the first case, no operator can wrap around the boundary since $y_{1}=0$ or $y_{n}=0$. We can thus impose obliviousness on $n-1$ qubits, where the Hamiltonian is arranged on a line. In the second case, we have $y_{1}=y_{n}=1$ so that an operator could wrap around the boundary. However, the above definition of obliviousness allows us to enumerate all possible ways to put tuples on the boundary. The usual requirement of obliviousness can then be imposed on the remaining qubits that are not on the boundary.

We now show that oblivious counting also results in a polynomial in $n$ in the periodic case. Throughout the remainder of this subsection, we use the notations

$$
\begin{align*}
C_{z y}(n, f):= & \left|\left\{\left\{i_{1}, \ldots, i_{\tau}\right\}: f\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=1, \operatorname{supp}\left(\left\{i_{1}, \ldots, i_{\tau}\right\}\right)=y, \operatorname{ival}(y)=z\right\}\right|  \tag{134}\\
C_{z y}\left(n, f_{j_{1}, \ldots, j_{\omega}}\right):=\mid\left\{\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}:\right. & f_{j_{1}, \ldots, j_{\omega}}\left(\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}\right)=1,  \tag{135}\\
& \left.\operatorname{supp}\left(\left\{k_{1}, \ldots, k_{\tau-\omega}\right\}\right)=\operatorname{infix}(y), \operatorname{ival}(y)=z\right\} \mid .
\end{align*}
$$

Theorem F.9. Let $G_{1}, \ldots, G_{\nu}$ be Hermitian operators and let $d_{1}, \ldots, d_{\nu} \in\{-1,1\}$. We assume that each term $G_{i}$ acts as identity except on the first $k$ qubits if $d_{i}=1$, or on the first $k-1$ qubits and the last qubit if $d_{i}=-1$. Suppose that $f: \mathcal{S}_{\tau}^{\mu} \rightarrow \mathbb{Z}_{2}$ induces an oblivious counting problem, as defined above. Then $\left|f^{-1}(1)\right|$ is a polynomial in $n$.
Proof. We compute $\left|f^{-1}(1)\right|$ by expanding it over all interval patterns and supports as

$$
\begin{align*}
\left|f^{-1}(1)\right| & =\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} \sum_{y \in \mathbb{Z}_{2}^{n}} C_{z y}(n, f)  \tag{136}\\
& =\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} \sum_{\substack{y \in \mathbb{Z}_{n}^{n} \\
y_{1} y_{n}=0}} C_{z y}(n, f)+\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} \sum_{\substack{y \in \mathbb{Z}_{n}^{n} \\
y_{1}=y_{n}=1}} C_{z y}(n, f) . \tag{137}
\end{align*}
$$

For the first term, no operator crosses the boundary. We can thus reduce the problem to counting $\tau$-tuples on $n-1$ qubits with interval pattern $z$ and support given by either the first or last $n-1$ positions of $y$. The Hamiltonian can effectively be regarded as arranged on a line. Invoking Theorem F.8, we see that the first term is a polynomial in $n$.

To handle the second term, we use a basic counting argument: to count the number of $\tau$-tuples, we add up all the possible cases where $\omega$ of them are on the boundary $\operatorname{bndry}(z)$ and the other $\tau-\omega$ have support infix $(y)$ away from the boundary. This allows us to rewrite

$$
\begin{equation*}
\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} \sum_{\substack{y \in \mathbb{Z}_{n}^{n} \\ y_{1}=y_{n}=1}} C_{z y}(n, f)=\sum_{z \in \mathbb{Z}_{\tau k}^{\tau}} \sum_{\substack{\left.\omega \in \llbracket \tau \rrbracket \\ \operatorname{supp}\left(\left\{j_{1}, \ldots, j_{\omega}\right\}\right\}\right)=\operatorname{bndry}(z)}} \sum_{\substack{\left\{j_{1} \\ y \in \mathbb{Z}_{1}^{n} \\ y_{1}=y_{n}=1\right.}} C_{z y}\left(n, f_{j_{1}, \ldots, j_{\omega}}\right) . \tag{138}
\end{equation*}
$$

By (O2) and a similar argument as above, the result is a sum of constant number of polynomials in $n$, which completes the proof.

Now we apply Theorem F. 9 to compute commutator bounds for our model Hamiltonian (1). To construct a product formula, we order the terms of the Hamiltonian similarly to the example in (114); in particular, we take them in the order

$$
\begin{equation*}
\sigma_{1}^{x} \sigma_{2}^{x}, \ldots, \sigma_{n-1}^{x} \sigma_{n}^{x}, \sigma_{n}^{x} \sigma_{1}^{x}, \sigma_{1}^{y} \sigma_{2}^{y}, \ldots, \sigma_{n-1}^{y} \sigma_{n}^{y}, \sigma_{n}^{y} \sigma_{1}^{y}, \sigma_{1}^{z} \sigma_{2}^{z}, \ldots, \sigma_{n-1}^{z} \sigma_{n}^{z}, \sigma_{n}^{z} \sigma_{1}^{z}, \sigma_{1}^{z}, \ldots, \sigma_{n}^{z} \tag{139}
\end{equation*}
$$

Our approach can be applied to product formulas of any order. However, we focus on the second- and fourth-order cases, for which the list of of operators in the product formulas have lengths $\mu=8 n$ and $\mu=40 n$, respectively. The resulting lists of terms take the form specified in Theorem F.9. Since each term in the Hamiltonian acts on at most two neighboring qubits, we can take the support of each operator to be a $0 / 1$ vector of Hamming weight 2, with a 1 in two consecutive positions.

To apply the commutator bounds, we must count $\tau$-tuples of Hamiltonian operators satisfying a certain commutation pattern, where $\tau \in\{2,3\}$ for the second-order bound and $\tau \in\{2,3,4,5\}$ for the fourth-order case. It remains to check that this counting problem is oblivious.

We first verify condition (O1). Fix an interval pattern $z \in \mathbb{Z}_{\tau k}^{\tau}$ and consider supports $y \in \mathbb{Z}_{2}^{n}$ and $y^{\prime} \in \mathbb{Z}_{2}^{n^{\prime}}$ with $\operatorname{ival}(y)=\operatorname{ival}\left(y^{\prime}\right)=z$ and $y_{1} y_{n}=y_{1}^{\prime} y_{n^{\prime}}^{\prime}=0$. By the latter conditions, we know that the relevant operators $H_{i_{1}}, \ldots, H_{i_{\tau}}$ do not touch the boundary. Intuitively, since the Hamiltonian includes all possible shifts of any given term, the counting problems for $y$ and $y^{\prime}$ should be equivalent.

To make this rigorous, we set up a one-to-one correspondence between $\tau$-tuples with support $y$ and those with support $y^{\prime}$. First consider the special case in which there is only one interval with the same length for both $y$ and $y^{\prime}$. Let $H_{j_{1}}, \ldots, H_{j_{\omega}}$ be operators with support supp $\left(\left\{j_{1}, \ldots, j_{\omega}\right\}\right)=y$. We construct a corresponding $\omega$-tuple with support $y^{\prime}$ as follows. Let $\delta$ be the position difference between the first non-zero element of $y^{\prime}$ and that of $y$. For any term in the Hamiltonian of the form $R^{d_{i} l} G_{i} R^{-d_{i} l}$, the corresponding operator is $R^{d_{i} l+\delta} G_{i} R^{-d_{i} l-\delta}$. More generally, we can apply the same transformation to all intervals in $y$ and $y^{\prime}$ with lengths $z_{1}, \ldots, z_{\tau}$ for any $\omega \in \llbracket \tau \rrbracket$ and $\left\{j_{1}, \ldots, j_{\omega}\right\} \in \mathcal{S}_{\omega}^{\mu}$. The resulting correspondence is easily checked to be invertible, so it specifies a bijection between $\omega$-tuples as claimed. We have thus established condition (O1). Condition (O2) follows similarly as $\operatorname{bndry}(z)$ and infix $(y)$ do not overlap.

Using symbolic polynomial interpolation, we derive the following succinct commutator bounds for the second- and fourth-order formulas.

Theorem F. 10 (Second-order commutator bound, succinct form). Let $H$ be the Hamiltonian (1), with terms ordered as in (139). Then the error in the second-order product formula approximation
satisfies

$$
\begin{equation*}
\left\|\exp (-i H t)-\left[S_{2}(-i t / r)\right]^{r}\right\| \leq \frac{\Lambda^{3}|t|^{3}}{r^{2}} T_{2}(n)+\frac{4(4 n \Lambda t)^{4}}{3 r^{3}} \exp \left(\frac{8 n \Lambda|t|}{r}\right), \tag{140}
\end{equation*}
$$

where

$$
T_{2}(n):= \begin{cases}194, & n=3  \tag{141}\\ 40 n^{2}-58 n, & n \geq 4\end{cases}
$$

Theorem F. 11 (Fourth-order commutator bound, succinct form). Let $H$ be the Hamiltonian (1), with terms ordered as in (139), and let $p:=1 /\left(4-4^{1 / 3}\right)$. Then the error in the fourth-order product formula approximation satisfies

$$
\begin{align*}
& \left\|\exp (-i H t)-\left[S_{4}(-i t / r)\right]^{r}\right\| \\
& \quad \leq\left(\frac{4 p-1}{2} \Lambda|t|\right)^{5} \frac{1}{5!\cdot r^{4}} T_{4}(n)+2 \frac{(20(4 p-1) n \Lambda t)^{6}}{6!\cdot r^{5}} \exp \left(\frac{20(4 p-1) n \Lambda|t|}{r}\right) \tag{142}
\end{align*}
$$

where

$$
T_{4}(n):= \begin{cases}23073564672, & n=3  \tag{143}\\ 94192316416, & n=4 \\ 278878851840, & n=5 \\ 1280000000 n^{4}-7701760000 n^{3}+23685120000 n^{2}-30224677632 n, & n \geq 6\end{cases}
$$

We conclude this appendix by considering the asymptotic gate complexity of the PF algorithm using our commutator bounds. Take the fourth-order bound as an example. With $\Lambda=\Theta(1)$ and $t=\Theta(n)$, the commutator bound is

$$
\begin{equation*}
\left\|\exp (-i H t)-\left[S_{4}(-i t / r)\right]^{r}\right\| \leq O\left(\frac{n^{9}}{r^{4}}+\frac{n^{12}}{r^{5}}\right) \tag{144}
\end{equation*}
$$

To guarantee that the simulation error $\epsilon$ is at most some constant, it suffices to use $r=O\left(n^{2.4}\right)$ segments. Since the circuit for each segment has size $O(n)$, we have an overall gate complexity of $O\left(n^{3.4}\right)$. Along similar lines, we find gate complexities of $O\left(n^{4}\right)$ (resp., $O\left(n^{11 / 3}\right)$ ) for the first-order (resp., second-order) commutator bound. These bounds improve the asymptotic gate complexities of the PF algorithm established by the analytic and minimized bounds (as shown in Table 1), which give $O\left(n^{5}\right)$ for the first-order formula, $O\left(n^{4}\right)$ for second order, and $O\left(n^{3.5}\right)$ for fourth order.

We only present concrete commutator bounds for the first-, second-, and fourth-order product formulas. In general, to obtain the $2 k$ th-order commutator bound, one must count the number of $(2 k+1)$-tuples satisfying a certain commutation pattern in a list of operators of length $8 \cdot 5^{k-1} n$. For $k \geq 3$, computing the exact form of the $(2 k)$ th order bound seems challenging even with the help of Theorem F. 9 and polynomial interpolation.

Nevertheless, it is still possible to obtain the asymptotic $n$-dependence of the commutator bound. The key step is to study those $(2 k+1)$-tuples for which all pairs of operators commute with each other. The number $N_{\text {comm }}$ of such commuting $(2 k+1)$-tuples satifsfies

$$
\begin{equation*}
\left(8 \cdot 5^{k-1}\right)^{2 k+1}\binom{n-2 k}{2 k+1} \leq N_{\mathrm{comm}} \leq\binom{ 8 \cdot 5^{k-1} n}{2 k+1} \tag{145}
\end{equation*}
$$

(Here the lower bound follows by placing each $(2 k+1)$-tuple so that its support contains exactly $2 k+1$ intervals.) Invoking Theorem F.9, we see that $N_{\text {comm }}$ is a polynomial in $n$ whose leading term is $\left(8 \cdot 5^{k-1} n\right)^{2 k+1}$.

When we Taylor expand the evolutions $\exp (\lambda H)$ and $S_{2 k}(\lambda)$, those $(2 k+1)$-tuples for which all pairs of operators commute with each other cancel. The remaining terms are either ( $2 k+1$ )-tuples where at least one pair of operators do not commute, or $l$-tuples with $l \leq 2 k$. Our above discussion shows that there are $O\left(n^{2 k}\right)$ such tuples. Therefore, the $2 k$ th-order commutator bound takes the form

$$
\begin{equation*}
\left\|\exp (-i H t)-\left[S_{2 k}(-i t / r)\right]^{r}\right\| \leq O\left(\frac{|t|^{2 k+1} n^{2 k}}{r^{2 k}}+\frac{(n t)^{2 k+2}}{r^{2 k+1}}\right)=O\left(\frac{n^{4 k+1}}{r^{2 k}}+\frac{n^{4 k+4}}{r^{2 k+1}}\right) \tag{146}
\end{equation*}
$$

To ensure that the simulation error is $O(1)$ for $t=n$, it suffices to choose $r=O\left(n^{2+2 /(2 k+1)}\right)$, which leads to a total gate complexity of $O\left(n^{3+2 /(2 k+1)}\right)$. This improves over the performance of the analytic and minimized bounds, which give complexity $O\left(n^{3+1 / k}\right)$.

## F. 3 Empirical bounds

While the bounds discussed in Section F. 1 and Section F. 2 provide rigorous correctness guarantees, they can be very loose. To understand the minimum resources that suffice for product formula simulation, we estimate their empirical performance. Of course, since quantum simulation is computationally challenging, we can only directly compute the actual simulation error for small instances. Using binary search, we find the value of $r$ (the total number of segments) that just suffices to ensure error $10^{-3}$. We extrapolate this behavior to produce a non-rigorous estimate of the performance of product formula simulation for instances of arbitrary size. We emphasize that the resulting empirical bound does not come with a guarantee of correctness. Nevertheless, we believe it better captures the true performance of product formula simulations and indicates the extent to which our rigorous bounds are loose.

We numerically simulate the product formula algorithm for systems of size 5 to 12 , determining the value of $r$ required to ensure error $10^{-3}$ as described above, and averaging over five random choices of the local field strengths $h_{j}$. We fit these data to power laws, as depicted in Figure 4. We find

$$
\begin{equation*}
r_{1}=2417 n^{1.964}, \quad r_{2}=39.47 n^{1.883}, \quad r_{4}=4.035 n^{1.555}, \quad r_{6}=1.789 n^{1.311}, \quad r_{8}=1.144 n^{1.141} \tag{147}
\end{equation*}
$$

where $r_{i}$ is the number of segments for the $i$ th-order formula to produce a simulation that is accurate to within $10^{-3}$. Considering the size of the circuit for each segment, this suggests an asymptotic complexity of roughly $9668 n^{2.964}$ for the first-order formula, $315.8 n^{2.883}$ for second order, $161.4 n^{2.555}$ for fourth order, $357.8 n^{2.311}$ for sixth order, and $1144 n^{2.141}$ for eighth order.

## G Taylor series implementation details

In this appendix, we discuss some technical details related to the implementation of the Taylor series (TS) algorithm, as introduced in Section C.2. We present concrete error bounds in Section G. 1 and discuss the failure probability of the approach in Section G.2. Then we describe how to implement select $(V)$ gates, a major component of the algorithm: in Section G.3, we explain how the control registers for these gates are encoded, and in Section G.4, we give an optimized implementation of $\operatorname{select}(V)$ gates.

## G. 1 Error analysis

In this section, we derive an explicit error bound that quantifies how the truncation order $K$ affects the overall accuracy. We then discuss how to apply this bound.


Figure 4: Comparison of the values of $r$ using the commutator and empirical bounds for formulas of order 1, 2 , and 4 , and values of $r$ for the empirical bound for formulas of order 6 and 8 . Straight lines show power-law fits to the data. The error bars for product formulas of order greater than 1 are negligibly small, so we omit them from the plots.

We begin by bounding the error introduced by truncating the Taylor series.
Lemma G.1. With the definitions of $\tilde{U}(t)$ in (7) and $t_{\mathrm{seg}}$ in (18), we have

$$
\begin{equation*}
\|\tilde{U}(t)-\exp (-i H t)\| \leq 2 \frac{(\ln 2)^{K+1}}{(K+1)!} \tag{148}
\end{equation*}
$$

for any $t \leq t_{\text {seg }}$.
Proof. We have

$$
\begin{align*}
\|\tilde{U}(t)-\exp (-i H t)\| & =\left\|\sum_{k=K+1}^{\infty} \frac{(-i H t)^{k}}{k!}\right\|  \tag{149}\\
& \leq \sum_{k=K+1}^{\infty} \frac{(\|H\| t)^{k}}{k!}  \tag{150}\\
& \leq \sum_{k=K+1}^{\infty} \frac{\left(\left(\alpha_{1}+\cdots+\alpha_{L}\right) t\right)^{k}}{k!}  \tag{151}\\
& \leq \sum_{k=K+1}^{\infty} \frac{(\ln 2)^{k}}{k!}  \tag{152}\\
& \leq 2 \frac{(\ln 2)^{K+1}}{(K+1)!}, \tag{153}
\end{align*}
$$

where the final inequality follows from Lemma F.2.
Next we consider the error induced by the isometry $\mathcal{V}(t)$ defined in (16). It is straightforward to verify that

$$
\begin{equation*}
(\langle 0| \otimes I) \mathcal{V}(t)=\frac{3}{s} \tilde{U}(t)-\frac{4}{s^{3}} \tilde{U}(t) \tilde{U}(t)^{\dagger} \tilde{U}(t) . \tag{154}
\end{equation*}
$$

As discussed in Section C.2, we rotate an ancilla qubit to increase the value of $s$ to be precisely $s=2$. Then the following bound characterizes the error in the implementation of the Taylor series.
Lemma G.2. Suppose $\|\tilde{U}-U\| \leq \delta$ for some unitary operator $U$. Then

$$
\begin{equation*}
\left\|\frac{3}{2} \tilde{U}-\frac{1}{2} \tilde{U} \tilde{U}^{\dagger} \tilde{U}-U\right\| \leq \frac{\delta^{2}+3 \delta+4}{2} \delta . \tag{155}
\end{equation*}
$$

Proof. Since $\|\tilde{U}-U\| \leq \delta$, we have

$$
\begin{equation*}
\|\tilde{U}\| \leq\|\tilde{U}-U\|+\|U\| \leq 1+\delta \tag{156}
\end{equation*}
$$

and therefore

$$
\begin{equation*}
\left\|\tilde{U} \tilde{U}^{\dagger}-I\right\| \leq\left\|(\tilde{U}-U) \tilde{U}^{\dagger}\right\|+\left\|U\left(\tilde{U}^{\dagger}-U^{\dagger}\right)\right\| \leq \delta(2+\delta) . \tag{157}
\end{equation*}
$$

Thus, by the triangle inequality, we have

$$
\begin{align*}
\left\|\frac{3}{2} \tilde{U}-\frac{1}{2} \tilde{U} \tilde{U}^{\dagger} \tilde{U}-U\right\| & \leq\|\tilde{U}-U\|+\frac{1}{2}\left\|\tilde{U}-\tilde{U} \tilde{U}^{\dagger} \tilde{U}\right\|  \tag{158}\\
& \leq \delta+\frac{\delta(2+\delta)(1+\delta)}{2}  \tag{159}\\
& =\frac{\delta^{2}+3 \delta+4}{2} \delta \tag{160}
\end{align*}
$$

as claimed.

We use the following basic property of contractions (operators of norm at most 1), which is easily proved using the triangle inequality.

Lemma G.3. Suppose $U_{i}: \mathcal{H} \rightarrow \mathcal{H}$ and $V_{i}: \mathcal{H} \rightarrow \mathcal{H}$ are contractions for all $i \in\{1, \ldots, r\}$. If $\left\|U_{i}-V_{i}\right\| \leq \eta$ for all $i$, then

$$
\begin{equation*}
\left\|U_{r} \cdots U_{2} U_{1}-V_{r} \cdots V_{2} V_{1}\right\| \leq r \eta . \tag{161}
\end{equation*}
$$

We also use the following lemma, which bounds the error introduced by normalization.
Lemma G.4. Suppose $\||\phi\rangle\|=1\|,|\psi\rangle \| \leq 1$, and $\||\psi\rangle-|\phi\rangle \| \leq \xi<1$. Then

$$
\begin{equation*}
\| \frac{|\psi\rangle}{\||\psi\rangle \|}-|\phi\rangle \| \leq \sqrt{1+\xi}-\sqrt{1-\xi} . \tag{162}
\end{equation*}
$$

Proof. Decompose $|\psi\rangle$ as

$$
\begin{equation*}
|\psi\rangle=\alpha|\phi\rangle+\beta\left|\phi^{\perp}\right\rangle \tag{163}
\end{equation*}
$$

for some normalized state $\left|\phi^{\perp}\right\rangle$ orthogonal to $|\phi\rangle$. Clearly $|\alpha|^{2}+|\beta|^{2} \leq 1$ since $\||\psi\rangle \| \leq 1$. Furthermore, the assumption $\||\psi\rangle-|\phi\rangle \| \leq \xi$ implies

$$
\begin{equation*}
|\alpha-1|^{2}+|\beta|^{2} \leq \xi^{2} \tag{164}
\end{equation*}
$$

so

$$
\begin{equation*}
|\alpha|^{2}+|\beta|^{2} \leq \xi^{2}+2 \operatorname{Re}(\alpha)-1 \tag{165}
\end{equation*}
$$

with $1-\xi \leq \operatorname{Re}(\alpha) \leq 1+\xi$.
Then we have

$$
\begin{align*}
\| \frac{|\psi\rangle}{\||\psi\rangle \|}-|\phi\rangle \| & =\sqrt{2-\frac{2 \operatorname{Re}(\alpha)}{\sqrt{|\alpha|^{2}+|\beta|^{2}}}}  \tag{166}\\
& \leq \sqrt{2-\frac{2 \operatorname{Re}(\alpha)}{\sqrt{\xi^{2}+2 \operatorname{Re}(\alpha)-1}}}  \tag{167}\\
& \leq \sqrt{2-\frac{2\left(1-\xi^{2}\right)}{\sqrt{\xi^{2}+2\left(1-\xi^{2}\right)-1}}}  \tag{168}\\
& =\sqrt{1+\xi}-\sqrt{1-\xi .} \tag{169}
\end{align*}
$$

Here the first inequality uses (165) and the fact that $\operatorname{Re}(\alpha) \geq 1-\xi \geq 0$, and the second inequality follows since the function $x / \sqrt{2 x-1+\xi^{2}}$ achieves its minimum at $x=1-\xi^{2}$ within the interval $1-\xi \leq x \leq 1+\xi$.

With all the above lemmas in hand, we are ready to prove an explicit error bound for the TS algorithm.

Theorem G.5. The TS algorithm achieves error at most $\sqrt{1+\xi}-\sqrt{1-\xi}$ with success probability at least $(1-\xi)^{2}$, where

$$
\begin{equation*}
\xi=r \frac{\delta^{2}+3 \delta+4}{2} \delta \quad \text { with } \quad \delta=2 \frac{(\ln 2)^{K+1}}{(K+1)!} \tag{170}
\end{equation*}
$$

| $K$ | $\xi / r$ |
| :---: | :---: |
| 1 | 1.3626 |
| 2 | 0.24118 |
| 3 | 0.039031 |
| 4 | 0.0053441 |
| 5 | 0.00061628 |
| 6 | $6.10123 \times 10^{-5}$ |
| 7 | $5.28621 \times 10^{-6}$ |
| 8 | $4.07124 \times 10^{-7}$ |
| 9 | $2.821965 \times 10^{-8}$ |
| 10 | $1.778215 \times 10^{-9}$ |

Table 2: Lookup table for the truncation order $K$, with $s$ boosted to be 2 in each segment.
Proof. For $t \in\left\{t_{\mathrm{seg}}, t_{\mathrm{rem}}\right\}$, Lemma G. 1 shows that

$$
\begin{equation*}
\|\tilde{U}(t)-\exp (-i H t)\| \leq \delta \tag{171}
\end{equation*}
$$

and Lemma G. 2 shows that

$$
\begin{equation*}
\|(\langle 0| \otimes I) \mathcal{V}(t)-\exp (-i H t) \| \leq \xi / r \tag{172}
\end{equation*}
$$

Since $\mathcal{V}(t)$ is an isometry, $(\langle 0| \otimes I) \mathcal{V}(t)$ is a contraction, so by Lemma G.3, we have

$$
\begin{equation*}
\|(\langle 0| \otimes I) \mathcal{V}\left(t_{\mathrm{rem}}\right)\left((\langle 0| \otimes I) \mathcal{V}\left(t_{\mathrm{seg}}\right)\right)^{r-1}-\exp (-i t H) \| \leq \xi \tag{173}
\end{equation*}
$$

The claim about the success probability follows by applying the triangle inequality, and the accuracy can be established by invoking Lemma G.4.

To apply this bound, we must determine the truncation order $K$ that achieves the desired error bound $\epsilon$. Just as for the product formula error bounds presented in Appendix F, it does not seem possible to compute $K$ in closed form. However, since $K$ can only take integer values, it is straightforward to tabulate the error estimates corresponding to all potentially relevant values of $K$, as shown in Table 2. Using the known value of $r$, we can then determine which value of $K$ suffices to ensure small error.

In Section F.3, we presented empirical error bounds for simulations based on product formulas. It would be natural to perform a similar analysis of the error in the TS algorithm. Unfortunately, it is intractable to find an empirical bound by direct simulation since the number of ancilla qubits used by the TS algorithm puts it beyond the reach of classical simulation even for very small systems (see Figure 2). A more limited alternative would be to use empirical data to improve the estimated error of the truncated Taylor series. However, since

$$
\begin{equation*}
\frac{(\ln 2)^{K+1}}{(K+1)!} \leq \sum_{k=K+1}^{\infty} \frac{(\ln 2)^{k}}{k!} \leq 2 \frac{(\ln 2)^{K+1}}{(K+1)!} \tag{174}
\end{equation*}
$$

the estimated error can be improved by a factor of at most 2 , which results in an additive offset of at most $\ln 2$ for the truncation order $K$. Thus we do not consider such a bound in our analysis.

## G. 2 Failure probability

Note that unlike the PF approach, the TS algorithm does not always succeed. When a measurement of the ancilla register indicates a failure, we must restart the simulation. To make a fair comparison between the PF and TS algorithms, we should take the success probability into account.

Fortunately, it turns out that the resulting overhead is almost negligible. With $\epsilon=10^{-3}$, Theorem G. 5 gives

$$
\begin{equation*}
\xi \leq \sqrt{\epsilon^{2}-\frac{\epsilon^{4}}{4}}<0.001 \tag{175}
\end{equation*}
$$

so the success probability of the algorithm is at least

$$
\begin{equation*}
(1-\xi)^{2} \geq 0.998 \tag{176}
\end{equation*}
$$

The expected number of times we must repeat the algorithm before succeeding is approximately $1 / 0.998 \approx 1.002$. Since this factor is very close to 1 , we simply neglect it, and we find it reasonable to directly compare gate counts for the PF algorithm (which has no probability of failure) and the TS approach.

## G. 3 Encoding of the control register

A crucial step in the implementation of the TS algorithm (and in the QSP algorithm) is to synthesize the select $(V)$ gates. The cost of this implementation depends strongly on the chosen representation for the control register.

Recall from (14) that the select $(V)$ operation has the form

$$
\begin{equation*}
\operatorname{select}(V):=\sum_{j=0}^{\Gamma-1}|j\rangle\langle j| \otimes V_{j} . \tag{177}
\end{equation*}
$$

For the TS algorithm, the operators $V_{j}$ are defined via (10). Here the index $j$ labels a value $k \in\{0,1, \ldots, K\}$ and indices $\ell_{1}, \ldots, \ell_{k} \in\{1, \ldots, L\}$. Perhaps the most straightforward approach is to represent the entire control register with a binary encoding using $\log _{2}(K+1)+K \log _{2} L$ bits. However, as pointed out in [14], we can significantly reduce the gate complexity by choosing a different encoding of the control register.

Specifically, we use a unary encoding to label $k$ and a binary encoding for each $\ell_{1}, \ldots, \ell_{k}$. With such an encoding, the instance of $\operatorname{select}(V)$ in the TS algorithm can be represented as the map

$$
\begin{equation*}
\left|1^{k} 0^{K-k}\right\rangle\left|\ell_{1}, \ldots, \ell_{K}\right\rangle|\psi\rangle \mapsto\left|1^{k} 0^{K-k}\right\rangle\left|\ell_{1}, \ldots, \ell_{K}\right\rangle(-i)^{k} H_{\ell_{1}} \cdots H_{\ell_{k}}|\psi\rangle . \tag{178}
\end{equation*}
$$

We implement this transformation as follows. Conditioned on the $j$ th qubit of the unary encoding of $k$ being 1 , and the $j$ th coordinate of $\ell_{1}, \ldots, \ell_{K}$ being the binary encoding of $\ell_{j}$, we apply $(-i) H_{\ell_{j}}$. Compared to an entirely binary encoding, this approach only requires an additional $\left\lceil K+1-\log _{2}(K+1)\right\rceil$ qubits, which is a modest increase since $K$ is typically small (see Table 2). In return, instead of selecting on a large register of $\Theta(K \log L)$ bits, we can perform $K+1$ independent selections on registers of $\log _{2} L$ bits, each controlled by a single qubit.

## G. 4 Implementation of select ( $V$ )

In this section, we present a circuit for the select $(V)$ operation $\sum_{\gamma=1}^{\Gamma}|\gamma\rangle\langle\gamma| \otimes V_{\gamma}$. We assume that the control register $\gamma \in\{1, \ldots, \Gamma\}$ is stored in binary using $w:=\left\lceil\log _{2} \Gamma\right\rceil$ bits. As discussed in

Section G.3, this procedure is a basic subroutine in our implementation of the select $(V)$ operation appearing in the TS algorithm, using a mixed unary/binary encoding. We also use this select $(V)$ construction in our implementation of the QSP algorithm.

Our goal is to apply a unitary operation $V_{\gamma}$ conditioned on the control register being in the state $|\gamma\rangle$. To do this, we cycle the value of a designated ancilla qubit through $\Gamma$ Boolean products of $w$ literals, where in each of the products, each of the variables $x_{1}, \ldots, x_{w}$ appears exactly once (either negated or not). We apply $V_{\gamma}$ conditioned on the ancilla qubit at the $\gamma$ th step of this construction. To simplify the presentation, we describe how to cycle through all $2^{w}$ Boolean products. This procedure can be terminated after producing only $\Gamma$ Boolean products to avoid unnecessary cost.

A straightforward way of obtaining the Boolean products is to implement them via multiplycontrolled Toffoli gates with appropriate control negations. Implementations of the multiplycontrolled Toffoli gates have been well-studied. In particular, using the approach of [55], and assuming that the negative controls introduce no additional cost, ${ }^{3}$ we can cycle through all Boolean products using the following resources:

$$
\begin{array}{rl}
4 w 2^{w}-6 \cdot 2^{w} & \text { Hadamard gates } \\
8 w 2^{w}-9 \cdot 2^{w} & T / T^{\dagger} \text { gates } \\
6 w 2^{w}-6 \cdot 2^{w} & \text { CNOT gates } \\
\left\lceil\frac{w-2}{2}\right\rceil & \text { ancilla qubits, initialized in and returned to the state }|0\rangle .
\end{array}
$$

We now present an optimized circuit to cycle through these Boolean products. In particular, we do this using a reversible circuit consisting of not, cnot, and Toffoli gates (we call a circuit over this gate set an $N C T$ circuit). Of the three basic gates used by the NCT circuits, the Toffoli gate is the most expensive - its best known implementation uses 6 CNOT gates, 2 Hadamard gates, and $7 T / T^{\dagger}$ gates [2]-so we first focus on optimizing the Toffoli gate count. We use Toffoli gates with controls of arbitrary polarities. All such Toffoli gates can be implemented by a Clifford $+T$ circuit using the same resources [2].

To cycle though $2^{w}$ Boolean products, clearly we must use at least $2^{w}$ gates (note that we do not allow gates with multiple targets). Indeed, most of these must be Toffoli gates, the most expensive gates in an NCT circuit.

Lemma G. 6 ([56, Lemma 4]). An NCT circuit for $\operatorname{select}(V)$ control generation requires at least $2^{w}-w-1$ Toffoli gates.

Considering the cost of a single Toffoli gate, a Clifford $+T$ circuit with $2^{w}$ Toffoli gates could contain as many as $6 \cdot 2^{w}$ CNOT gates and $7 \cdot 2^{w} T / T^{\dagger}$ gates. We next construct a quantum circuit that cycles through all $2^{w}$ Boolean products, and show that after several optimizations, the gate counts are close to the above optimistic figure, and significantly better than those obtained using multiply-controlled Toffoli gates. We begin by describing an efficient NCT circuit that cycles through $2^{w}$ Boolean products.

Lemma G.7. The select $(V)$ control generation can be implemented by an NCT circuit using at most 2 nOt gates, $1.5 \cdot 2^{w}-2$ CNOT gates, and $1.5 \cdot 2^{w}-4$ Toffoli gates, with $w-1$ ancilla qubits initialized in and returned to the state $|0\rangle$.

Proof. Consider a depth- $w$ complete binary tree, where vertices are labeled by the Boolean products of the $w$ variables $x_{1}, \ldots, x_{w}$. A vertex at depth $d$ is labeled by a product of the first $d$ variables,

[^2]

Figure 5: Binary tree encoding the circuit that implements the control generation for the select $(V)$ operation.
$x_{1}^{a_{1}} \ldots x_{d}^{a_{d}}$, where $x_{i}^{a_{i}}:=x_{i}$ if $a_{i}=$ LEFT and $x_{i}^{a_{i}}:=\overline{x_{i}}$ if $a_{i}=$ RIGHT. The label $a_{d}$ indicates how we arrive at $x_{1}^{a_{1}} \ldots x_{d}^{a_{d}}$ from its parent $x_{1}^{a_{1}} \ldots x_{d-1}^{a_{d-1}}$, by taking the left edge if $a_{d}=$ LEFT or the right edge if $a_{d}=$ RIGHT. The product $x_{1}^{a_{1}} \ldots x_{d}^{a_{d}}$ thus indicates the path to the vertex from the root, as shown in Figure 5. Leaves of the tree are labeled by the Boolean products over all $w$ variables $x_{1}, \ldots, x_{w}$, in all possible polarities.

In the quantum circuit implementing the control generation for the select $(V)$ transformation, all vertices at a given depth are represented by a single qubit, labeled $q_{d}$ for $d \in\{0,1, \ldots, w\}$. The qubit $q_{k}$ carries the product of the first $k$ variables with appropriate polarities that evolve as the computation proceeds. We discuss how to implement the circuit using $w+1$ ancilla qubits $q_{0}, q_{1}, \ldots, q_{w}$, although it is straightforward to reduce the number of ancillas to $w-1$. The root of the tree corresponds to the Boolean constant 1, so it can be viewed as a virtual qubit. The two depth- 1 vertices represent the variable $x_{1}$ and its negation, values that can be obtained in-place by negating the variable $x_{1}$ as needed. Thus it suffices to use $w-1$ ancilla qubits $q_{2}, q_{3}, \ldots, q_{w}$.

We compute all desired Boolean products by traversing the tree using the left-hand rule, i.e., always taking the leftmost edge going down if there are vertices not yet explored, and going up if all edges and vertices below were already explored. The path starts and ends in the root (marked with an asterisk in Figure 5), and traverses each edge twice, once going down and once coming up, so it has length $4 \cdot 2^{w}-4$.

When traversing the edge $\left(x_{1}^{a_{1}} x_{2}^{a_{2}} \ldots x_{d-1}^{a_{d-1}}, x_{1}^{a_{1}} x_{2}^{a_{2}} \ldots x_{d-1}^{a_{d-1}} x_{d}^{a_{d}}\right)$ in either direction, we apply the gate $\operatorname{TOF}\left(q_{d-1}, x_{d}^{a_{d}} ; q_{d}\right)$, where $\operatorname{TOF}\left(c_{1}, c_{2} ; t\right)$ denotes a Toffoli gate with controls $c_{1}$ and $c_{2}$ and target $t$. These edges in the walk are colored blue in Figure 5. No gate needs be applied for edges connecting to the root, as the variable $x_{1}$ and its negation can be obtained without using a Toffoli or a cnot gate. Those edges in the walk are colored grey in Figure 5. Overall, this circuit uses $4 \cdot 2^{w}-8$ Toffoli gates and two not gates (one to negate $x_{1}$ and the other to uncompute this negation). During the walk, each time we reach a leaf, we create a new Boolean product of $w$ variables with polarities determined by the path to the root. By visiting all leaves, we create all $2^{w}$ desired products. Completing the path and returning to the root ensures that all ancilla qubits are reset to 0 . Note that we can return to the root from any given leaf without visiting all leaves, finishing the walk early if only some of the Boolean products are needed (and thereby using fewer gates).


Figure 6: Optimized circuit walk on the binary tree, using grey, blue, red, and green solid edges. Dashed edges indicate reductions in the quantum gate count.

Next we show how to reduce the Toffoli gate count. Observe that the above path makes $2^{w}-2$ moves of the form $v_{1} \rightarrow v_{2} \rightarrow v_{3}$, where $v_{1}$ and $v_{3}$ are the children of $v_{2}$, which is not the root (for example, consider any path indicated by a red arc in Figure 6). Such a path requires applying the gates $\operatorname{TOF}\left(q_{d-1}, x_{d}^{\mathrm{LEFT}} ; q_{d}\right)$ and $\operatorname{TOF}\left(q_{d-1}, x_{d}^{\mathrm{RIGHT}} ; q_{d}\right)$. However, by the circuit identity

$$
\begin{equation*}
\operatorname{TOF}\left(q_{d-1}, x_{d}^{\mathrm{LEFT}} ; q_{d}\right) \operatorname{TOF}\left(q_{d-1}, x_{d}^{\mathrm{RIGHT}} ; q_{d}\right)=\operatorname{CNOT}\left(q_{d-1} ; q_{d}\right) \tag{179}
\end{equation*}
$$

we can replace two Toffoli gates with one CNOT gate. Since $v_{2}$ is never a leaf, we do not need to explicitly visit it, thereby saving resources. This modification reduces the Toffoli gate count by $2\left(2^{w}-2\right)=2 \cdot 2^{w}-4$, giving a reduced Toffoli gate count of $2 \cdot 2^{w}-4$. However, there are now $2^{w}-2$ CnOTs. The circuit still uses 2 not gates.

Next, consider a path $v_{1} \rightarrow v_{2} \rightarrow v_{3} \rightarrow v_{4} \rightarrow v_{5}$, where $v_{2}$ and $v_{4}$ are the children of $v_{3}, v_{1}$ is the right child of $v_{2}$, and $v_{5}$ is the left child of $v_{4}$. Such paths are indicated by green arcs in Figure 6 . Taking such a path involves applying the circuit

$$
\begin{equation*}
\operatorname{TOF}\left(q_{d-1}, \overline{x_{d}} ; q_{d}\right) \operatorname{CNOT}\left(q_{d-2} ; q_{d-1}\right) \operatorname{TOF}\left(q_{d-1}, x_{d} ; q_{d}\right) \tag{180}
\end{equation*}
$$

This gate sequence can be simplified to

$$
\begin{equation*}
\operatorname{CNOT}\left(q_{d-1} ; q_{d}\right) \operatorname{TOF}\left(q_{d-2}, x_{d} ; q_{d}\right) \operatorname{CNOT}\left(q_{d-2} ; q_{d-1}\right) \tag{181}
\end{equation*}
$$

This additional optimization eliminates $\frac{1}{2} \cdot 2^{w}-1$ Toffoli gates and introduces $\frac{1}{2} \cdot 2^{w}-1$ CNOTs, giving a reduced Toffoli gate count of $2 \cdot 2^{w}-4-\frac{1}{2} \cdot 2^{w}+1=1.5 \cdot 2^{w}-3$.

We conclude by carefully counting the resources, while reducing them a little further. Naively, every move along an edge of the tree involves applying a Toffoli gate. As detailed above, each red shortcut lets us traverse two edges using only one CNOT gate and each green shortcut lets us traverse four edges using one Toffoli and two cnots. Overall, the number of Toffoli gates, per the
colored tree in Figure 6, is

$$
\begin{array}{ll}
4 \cdot 2^{w}-4 & \\
-2^{w} & \\
-1.5 \cdot 2^{w}+3 & \\
\text { (length of the entire unreduced path) } \\
-2 & \text { (green shortcuts reduction) } \\
-1 & \text { (grey edges connect to a virtual node and require no Toffoli gates) }  \tag{182}\\
=1.5 \cdot 2^{w}-4 . &
\end{array}
$$

The cnot gate count is

$$
\begin{align*}
0.5 & \cdot 2^{w} & & \text { (red) } \\
& +2 \cdot\left(2^{w-1}-1\right) & & \text { (green) } \\
& +0 & & \text { (across blue and grey paths) } \\
= & 1.5 \cdot 2^{w}-2 . & & \tag{183}
\end{align*}
$$

The not gate count is 2 .
By substituting Toffoli gates with their Clifford $+T$ implementations, the desired control generation circuit may be implemented using 2 not gates, $3 \cdot 2^{w}-8$ Hadamard gates, $10.5 \cdot 2^{w}-28 T$ gates, and $10.5 \cdot 2^{w}-26$ CNOT gates. In the remainder of this section, we reduce this gate count. To accomplish this, we pair Toffoli gates as indicated in Figure 6 with dashed blue lines. Formally, we pair every $\operatorname{TOF}\left(q_{d-1}, \overline{x_{d}} ; q_{d}\right)$ with the $\operatorname{TOF}\left(q_{d-1}, x_{d} ; q_{d}\right)$, such that the tree walk path between these two edges does not visit any leaf, and the gates applied between these two Toffoli gates do not affect qubits $x_{d}$ and $q_{d}$. The path beginning and ending with such a pair of Toffoli gates results in the following circuit, where the unitary $R$ implements all gates in between:


The two Toffoli gates can be decomposed into Clifford $+T$ gates such that the $\operatorname{TOF}\left(q_{d-1}, \overline{x_{d}} ; q_{d}\right)$ gate ends with one Hadamard, two cnots, and $3 T / T^{\dagger}$ gates on qubits $x_{d}$ and $q_{d}$. These cancel with one Hadamard, two CNOTs, and $3 T / T^{\dagger}$ gates in a decomposition of $\operatorname{TOF}\left(q_{d-1}, x_{d} ; q_{d}\right)$, leaving behind a single Phase gate (due to the Toffoli gate controls appearing with different polarities). We furthermore choose the minimal resource implementation of $\operatorname{TOF}\left(q_{d-1}, \overline{x_{d}} ; q_{d}\right)$ such that the last gate applying to qubit $\left|q_{d-1}\right\rangle$ is $\operatorname{CNOT}\left(q_{d} ; q_{d-1}\right)$ and the first gate in the decomposition of $\operatorname{TOF}\left(q_{d-1}, x_{d} ; q_{d}\right)$ affecting qubit $\left|q_{d-1}\right\rangle$ is also CNOT $\left(q_{d} ; q_{d-1}\right)$. These two CNOTs can be commuted to each other and canceled, since Phase commutes with the control on the qubit $\left|q_{d}\right\rangle$ and $R$ is a reversible circuit with its output xored onto the target qubit $\left|q_{d-1}\right\rangle$, implying that $R$ commutes with $\operatorname{CNOT}\left(q_{d} ; q_{d-1}\right)$. The reduced circuit, equivalent to the one shown above, is as follows:


The equivalence of the circuits before and after optimization can be verified by direct computation.
The control generation circuit contains $\frac{1}{2} \cdot 2^{w}-w$ such Toffoli pairs, each reducing the gate count by 2 Hadamards, $6 T / T^{\dagger}$ gates, and 6 cnots while introducing a single Phase gate. Overall, the optimized select $(V)$ control generation circuit thus uses the following resources:

$$
\begin{array}{rl}
2 & \text { NOT gates } \\
2 \cdot 2^{w}+2 w-8 & \text { Hadamard gates } \\
0.5 \cdot 2^{w}-w & \text { Phase gates } \\
7.5 \cdot 2^{w}+6 w-28 & T / T^{\dagger} \text { gates } \\
7.5 \cdot 2^{w}+6 w-26 & \text { CNOT gates } \\
w-1 & \text { ancilla qubits, initialized in and returned to the state }|0\rangle .
\end{array}
$$

Note that the resulting $T / T^{\dagger}$-count is only about $7 \%$ higher than the expected $T / T^{\dagger}$-count in a circuit with $2^{w}-w-1$ Toffoli gates (which is the number of Toffoli gates required to generate $2^{w}$ Boolean products by a reversible NCT circuit, according to Lemma G.6). The overall elementary gate count in our optimized implementation is $17.5 \cdot 2^{w}+o\left(2^{w}\right)$. This compares favorably to the baseline estimation of $18 w 2^{w}+o\left(w 2^{w}\right)$ elementary gates, improving the gate count by a factor of $\frac{36 w}{35}$ at leading order. Recall that in our application to the TS algorithm, we have $w=\left\lceil\log _{2} \Gamma\right\rceil$, where $\Gamma$ is the number of selected unitaries, as in (177).

## H Quantum signal processing implementation details

We now turn our attention to the quantum signal processing (QSP) algorithm of Low and Chuang $[51,52]$, as introduced in Section C.3. We estimate the success probability of the QSP algorithm in Section H. 1 and discuss how this can be taken into account to make a fair comparison with other simulation algorithms. In Section H.2, we describe optimizations that reduce the gate count of our implementation. Then, in Section H.3, we discuss the difficulty of computing the phases that specify the algorithm and describe a segmented version of the algorithm that mitigates this issue. Finally, in Section H.4, we describe empirical bounds on the error in the truncated Jacobi-Anger expansion and in the overall algorithm.

## H. 1 Failure probability

Just like the TS algorithm, the QSP algorithm is probabilistic. The success probability is not explicitly evaluated in references [51,52]. The error analysis presented in Section C. 3 guarantees that the post-selected operator $(\langle+|\langle G| \otimes I) V(|+\rangle|G\rangle \otimes I)$ is close to the ideal evolution $e^{-i H t}$ in the sense that

$$
\begin{equation*}
\left\|(\langle+|\langle G| \otimes I) V(|+\rangle|G\rangle \otimes I)-e^{-i H t}\right\| \leq \epsilon . \tag{184}
\end{equation*}
$$

Therefore, for any given input state $|\psi\rangle$, the postselection succeeds with probability at least

$$
\begin{align*}
& \|(\langle+|\langle G| \otimes I) V(|+\rangle|G\rangle \otimes|\psi\rangle) \|^{2} \\
& \quad \geq\left(\| e^{-i H t}|\psi\rangle\|-\|(\langle+|\langle G| \otimes I) V(|+\rangle|G\rangle \otimes|\psi\rangle)-e^{-i H t}|\psi\rangle \|\right)^{2}  \tag{185}\\
& \quad \geq\left(1-\left\|(\langle+|\langle G| \otimes I) V(|+\rangle|G\rangle \otimes I)-e^{-i H t}\right\|\right)^{2}  \tag{186}\\
& \quad \geq 1-2 \epsilon . \tag{187}
\end{align*}
$$

With $\epsilon=10^{-3}$, the success probability of the QSP algorithm is at least $1-2 \epsilon=0.998$. Reasoning similarly as in Section G.2, we find that the expected number of times we must repeat the algorithm before succeeding is approximately $1 / 0.998 \approx 1.002$. This factor is very close to 1 , so as for the TS algorithm, we neglect it when comparing gate counts.

## H. 2 Circuit optimizations

The select $(V)$ gate is a major component of the QSP algorithm, so we use the optimized implementation of that subroutine described in Section G. 4 in our implementation of the QSP algorithm. In this section, we present some further circuit optimizations that also reduce the gate count.

As discussed in Section C.3, we use the phased iterate $V_{\phi}$ defined in (31), whereas Low and Chuang use the operation $V_{\phi}^{\prime}$ defined in (40). It is easy to see that

$$
\begin{equation*}
V_{\phi}=\left(I \otimes Z_{\pi / 2}\right) V_{\phi}^{\prime}\left(I \otimes Z_{-\pi / 2}\right), \tag{188}
\end{equation*}
$$

so

$$
\begin{equation*}
V_{\phi}^{\dagger}=\left(I \otimes Z_{\pi / 2}\right) V_{\phi}^{\prime \dagger}\left(I \otimes Z_{-\pi / 2}\right) \tag{189}
\end{equation*}
$$

also involves conjugation by $I \otimes Z_{\pi / 2}$. Thus, when the phased iterates are applied in the sequence (35), the inner partial reflection gates cancel. Furthermore, $Z_{\varphi}$ simply introduces a relative phase between $|G\rangle$ and its orthogonal subspace, so its action is trivial if the ancilla register is initialized in and postselected on $|G\rangle$. Thus we see that the implementation as defined in Section C. 3 has the same effect as in [51]. Each partial reflection is implemented using $O(\log n)$ elementary gates, and there are $O\left(n^{2}\right)$ phased iterates, so our implementation saves $O\left(n^{2} \log n\right)$ gates.

We apply a similar simplification to further reduce the gate count. For every phased iterate $V_{\phi}$ defined in (31), we must implement a controlled version of the operator $-i Q=-i((2|G\rangle\langle G|-$ $I) \otimes I)$ select $(H)$. In particular, this requires us to perform a controlled-reflection about $|G\rangle$. We can do this by performing a controlled- $U^{\dagger}$ that unprepares the state $|G\rangle$ (where $U$ is a unitary operation satisfying $U|0\rangle=|G\rangle$ ), a controlled reflection about $|0\rangle$, and finally a controlled- $U$ that prepares the state $|G\rangle$. However, observe that we can replace the controlled unitary conjugation by its uncontrolled version without changing the behavior of the circuit. Furthermore, by grouping neighboring pairs of phased iterates in the sequence of $V_{\phi}$ and $V_{\phi^{\prime}}^{\dagger}$ operations, we can cancel pairs of unitary operators $U$ and $U^{\dagger}$ for state preparation and unpreparation.

## H. 3 Phase computation and segmented algorithm

Recall that to specify the QSP algorithm, we must find phases $\phi_{1}, \ldots, \phi_{M}$ that realize the truncated Jacobi-Anger expansion. In principle, these angles can be computed in polynomial time [53]. However, this computation is difficult in practice, so we can only carry it out for very small instances. Specifically, we found the time required to calculate the angles to be prohibitive for values of $M$ greater than about 32. For $n=10$ qubits with $t=n$ and $\epsilon=10^{-3}$, the error bound (38) suggests that we should take $M=1100$. Thus the difficulty of computing the angles prevents us from synthesizing nontrivial instances of the algorithm. This difficulty arises because the procedure for computing the angles requires us to compute the roots of a high-degree polynomial to high precision. It is a natural open problem to give a more practical method for computing the angles.

Fortunately, to determine the Clifford $+R_{z}$ gate count in our implementation of the QSP algorithm, we do not need to know the angles of the phased iterates. Furthermore, since most $R_{z}$ gates require approximately the same number of Hadamard and $T$ gates to realize within a given precision, we can get a reasonable estimate of the Clifford $+T$ count by using random angles in place of the true values. However, we emphasize that this method does not produce a correct quantum simulation circuit, and should only be used as a benchmark of the resource requirements of the QSP algorithm - which is only useful if the true angles can ultimately be computed.

An alternative is to consider what we call a segmented version of the algorithm. In this approach, we first divide the evolution time into $r$ segments, each of which is sufficiently short that the angles


Figure 7: Comparison of the number of phased iterates using optimal and segmented implementations of the QSP algorithm.
can readily be computed. Since the optimality of the QSP approach to Hamiltonian simulation relies essentially on simulating the entire evolution as a single segment, the segmented approach has higher asymptotic complexity. However, it allows us to give a complete implementation, and the overhead for moderate values of $n$ is not too great.

To analyze the algorithm with $r$ segments, we apply the error bound (38) with $t$ replaced by $t / r$ and $\epsilon$ replaced by $\epsilon / r$. This gives the sufficient condition

$$
\begin{equation*}
\frac{4(\alpha t / r)^{q}}{2^{q} q!} \leq \frac{\epsilon}{8 r} \tag{190}
\end{equation*}
$$

where $q=\frac{M}{2}+1$. Thus

$$
\begin{equation*}
r=O\left(\alpha t\left(\frac{\alpha t}{2 \epsilon\left(\frac{M}{2}+1\right)!}\right)^{2 / M}\right) \tag{191}
\end{equation*}
$$

segments suffice to ensure overall error at most $\epsilon$. With $t=n, \alpha=O(n)$, and $M$ a fixed constant, we have $r=O\left(n^{2+4 / M}\right)$ segments. Within each segment, the number of phased iterates is $M$, which is independent of the system size. Each phased iterate has circuit size $O(n)$ using the improved select $(V)$ implementation described in Section G.4. Therefore, the segmented algorithm has gate complexity $O\left(n^{3+4 / M}\right)$.

In our implementation, we use $M=28$ (i.e., $q=15$ ). For the instance of quantum simulation considered in this paper, we set $\epsilon=10^{-3}, \alpha=4 n$, and $t=n$. With these values, (190) shows that it suffices to use

$$
\begin{equation*}
r \geq \sqrt[14]{\frac{10^{3} \cdot 2^{20} n^{30}}{15!}}=0.6010 n^{15 / 7} \tag{192}
\end{equation*}
$$

segments.
Figure 7 compares the total number of phased iterates used in the segmented and optimal implementations. Over the range of interest, the segmented algorithm is only worse by a factor between 2 and 3 .


Figure 8: Comparison of the number of phased iterates using the analytic bound (37) and the empirical bound for the Jacobi-Anger expansion. Here $M$ is the number of phased iterates and $n$ is the system size.

## H. 4 Empirical error bounds

The error bound (38) uses the closed-form expression (37) for the remainder of the Jacobi-Anger expansion. While it is a convenient to use such an analytical expression, it is natural to ask how tightly it bounds the complexity of the algorithm.

To address this question, we numerically evaluate the left-hand side of (37) for systems of sizes ranging from 10 to 20, as shown in Figure 8. By extrapolating these data, we estimate the complexity of the QSP algorithm for arbitrary sizes, including those for which classical evaluation of the series is intractable. The empirical bound improves the gate count by a factor between 1.25 and 1.45 over the range of interest $(10 \leq n \leq 100)$. More specifically, power law fits to the data give

$$
\begin{equation*}
M_{\mathrm{ana}}=11.30 n^{1.988}, \quad M_{\mathrm{emp}}=9.849 n^{1.939} \tag{193}
\end{equation*}
$$

for the number of phased iterates using either the analytic bound or the empirical bound, respectively. Since each phased iterate has gate complexity $O(n)$ (using the technique from Section G.4), we find that the QSP algorithm has complexity $O\left(n^{2.988}\right)$ (resp., $O\left(n^{2.939}\right)$ ) using the analytic bound (resp., empirical bound). Note that the former is roughly consistent with the upper bound $O\left(n^{3} \log n\right)$ shown in Table 1.

We do not consider the empirical bound for the segmented version of the QSP algorithm, since the savings is small in that case (even less at $M=28$ than at the values shown in Figure 8), and the main goal of the segmented approach is to have a fully-specified algorithm with rigorous guarantees. However, we use the empirical bound to estimate resources using the non-segmented QSP algorithm. This produces our most optimistic benchmark for the performance of the QSP algorithm.

One could also consider a full empirical bound for the QSP algorithm by using direct simulation to determine its true overall error. The need for ancilla qubits makes this challenging: the algorithm uses $n+\lceil\log 4 n\rceil+1$ qubits to simulate an $n$-qubit system, as shown in Figure 2. Fortunately, unlike with the TS algorithm, small instances of the QSP algorithm are just within reach of direct classical


Figure 9: Empirical error in the segmented QSP algorithm and PF algorithms of orders 1, 2, and 4 (with commutator bound) for small system sizes.
simulation.
However, preliminary numerical investigation suggests that the performance of the QSP algorithm cannot be significantly improved using such an empirical bound. Figure 9 shows the empirical error in the segmented QSP algorithm for small system sizes, averaging over 10 random experiments with a fixed target error $\epsilon=10^{-3}$, along with similar data for the PF algorithm using the commutator bound. We observe that for system sizes between 5 and 9 , the QSP error is consistently around $5 \times 10^{-5}$, which is not significantly less than the target error of $10^{-3}$. While there was more variation in the error of the QSP algorithm as compared to the PF algorithm, in no case was the QSP error less than $10^{-5}$. In contrast to the PF algorithm, where the error apparently decreases as a power law in $n$, the QSP error shows no indication of decreasing. Furthermore, since the complexity of the QSP algorithm depends logarithmically on the inverse error $1 / \epsilon$, even a large reduction in the error may not have a significant effect. For these reasons, we do not consider full empirical error bounds in our resource estimates for the QSP algorithm.

## I Detailed results

In this appendix, we present our results in more detail, expanding upon the discussion in Section 4. First we briefly describe the variation of gate counts with respect to the random choice of

Hamiltonian (namely, the coefficients $h_{j} \in[-1,1]$ in (1)). For each of the system sizes considered below, we generated five random Hamiltonians and determined the gate counts using Quipper. We report the average gate count over the five Hamiltonians.

In the Clifford $+R_{z}$ basis, the gate counts for PF algorithms are actually independent of the particular choice of coefficients $h_{j}$ : the rotation angles in the circuits differ, but the number of rotations does not. The number of gates in the Clifford $+T$ basis depends on these angles, but only weakly. Across orders and bounds, the $T$ count never varies by more than $1 \%$. The TS algorithm is the most sensitive to the choice of Hamiltonian. For systems of size 13, the Clifford $+R_{z}$ gate counts vary by about $10 \%$, and similarly for the $T$ counts. However, the variance reduces as the system size increases. For systems of size 100, the variance is about $2 \%$ across gate types. The QSP algorithm (segmented or not) has a mild dependence on the random choice of Hamiltonian. In the Clifford $+R_{z}$ basis, the CNOT and $R_{z}$ counts are independent of the choice of Hamiltonian while the Clifford counts vary by less than $1 \%$. Because the rotation angles depend on the Hamiltonian, the $T$ count depends on the choice of Hamiltonian, but again the observed variation is always less than $1 \%$.

In the remainder of this section, we present our results in detail. We discuss the PF algorithm first, in Section I.1, since this case has the most variations to compare. Then, in Section I.2, we present results for the TS and QSP algorithms and discuss the relative performance of all three approaches. In Section I.3, we explain how the circuit counts are improved using our automated circuit optimizer.

## I. 1 Product formulas

The performance of quantum simulation algorithms using product formulas depends strongly on the order of the formula and the method used to bound its error. We now compare the effect of these choices.

Figure 10 compares the total gate counts of simulation algorithms using product formulas of various orders with each of our four error bounds. Using rigorous bounds, for system sizes between 13 and 500 , the fourth-order algorithm outperforms other orders except for very small system sizes. The second-order algorithm is preferred for the largest range of system sizes with the minimized bound, but even in that case, it only outperforms the fourth-order algorithm for $n \leq 28$. Similar considerations hold for Clifford $+T$ circuits and when we only count gates of a particular type. Overall, it is clear that the fourth-order algorithm using the commutator bound is preferred for our application if one requires a provable guarantee on the error. If a heuristic method is acceptable, the empirical bound offers a dramatic reduction in complexity. In that case, the sixth-order algorithm begins to outperform the fourth-order algorithm at around 30 qubits, although the difference is small for the range of sizes we focus on (at most 100 qubits).

Of the cases we considered, the first-order algorithm has the worst performance, using from $10^{2}$ to over $10^{4}$ times more gates than the fourth- or sixth-order algorithm. The experimental demonstrations of digital quantum simulation that we are aware of $[6,19,48]$ have primarily used the first-order formula, with some limited applications of the second-order formula [19, 48]. Our work suggests that higher-order formulas may be practically relevant for surprisingly small instances, and we hope our results motivate practitioners to take advantage of this.

We did not evaluate the commutator bound for orders higher than 4 . This is because the bound is difficult to compute in practice, as discussed in Section F.2. Since the fourth-order algorithm performs significantly better than the sixth-order one for both the analytic and minimized bounds, we also expect the fourth-order algorithm to outperform the sixth-order algorithm even if we used an analogous sixth-order commutator bound.


Figure 10: Total gate counts in the Clifford $+R_{z}$ basis for product formula algorithms using the analytic (top left), minimized (top right), commutator (bottom left), and empirical (bottom right) bounds, for system sizes between 13 and 500 .

Figure 11 shows the total gate counts for the Clifford $+R_{z}$ gate set using the fourth-order algorithm with the analytic, minimized, commutator, and empirical bounds. The commutator bound saves a factor of more than 10 over the minimized bound, a significant improvement in practical terms. The empirical bound results in even greater savings of a further factor of about 100 (for system sizes in the 10s) to over 500 (for system sizes around 100). Clearly, the empirical bound is strongly preferred if one can tolerate its lack of a rigorous correctness guarantee. This suggests that existing rigorous bounds are quite loose, so a natural open problem is to establish stronger rigorous bounds, as mentioned in Section 5 .

In fact, the commutator and empirical bounds not only reduce the gate count for small instances, but actually improve the dependence on system size $n$. This is evident in Figure 11, where the slope of the data is reduced for the commutator and empirical bounds. Table 3 presents the exponents of $n$, both for the proven performance of the commutator bound (as stated in Theorem F.5, Theorem F.10, and Theorem F.11) and the numerically observed performance of the empirical bound. Surprisingly, the empirical bound appears to give better asymptotic performance than $O\left(n^{3}\right)$ (the best asymptotic gate complexity listed in Table 1): in particular, the eighth-order formula has asymptotic empirical complexity only slightly greater than quadratic in $n$.

While our empirical bound is only directly applicable to the particular system and parameters


Figure 11: Total gate counts of Clifford $+R_{z}$ circuits using the fourth-order PF algorithm and varying bounds for system sizes between 10 and 500 .

| Bound | Order |  |  |  |  |  |  |  |  | $\theta$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 1 | 2 | 4 | 6 | 8 |  |  |  |  |  |  |  |
| - Analytic/Minimized | 5 | 4 | 3.5 | 3.333 | 3.25 |  |  | 0 |  |  |  | $\bullet$ |
| - Commutator | 4 | 3.667 | 3.4 | 3.286 | 3.222 |  |  |  |  |  |  | $\bigcirc$ |
| - Empirical | 2.964 | 2.883 | 2.555 | 2.311 | 2.141 |  |  | 1 | 2 |  | 6 | 8 |

Table 3: Exponent of $n$ in the asymptotic gate complexity of product formula simulation algorithms using various error bounds.
considered in our study, we expect that qualitatively similar improvements should apply in other cases. In particular, we performed a similar extrapolation for the fourth-order product formula applied to the same system with $h=1 / 10$ and $h=10$. We found very similar performance estimates for these parameters, with essentially the same exponent of $n$ and slightly different prefactors.

Note that while the constants in high-order commutator bounds are nontrivial to compute, we can more easily determine the asymptotic performance of these bounds, as discussed at the end of Section F.2. The lowest-order error in the $(2 k)$ th-order product formula is $O\left(n^{4 k+2} / r^{2 k+1}\right)$. Our commutator bound utilizes the commutation relations to reduce the $n$-dependence of this lowestorder error to $O\left(n^{4 k+1} / r^{2 k+1}\right)$, while the remaining higher-order terms are handled by standard techniques [11]. The resulting asymptotic performance is $O\left(n^{3+2 /(2 k+1)}\right)$, improving over the complexity $O\left(n^{3+1 / k}\right)$ obtained using the standard bound.

## I. 2 Other algorithms and comparisons

Figure 12 compares gate counts for the fourth-order PF algorithm with commutator bound, the better of the fourth- or sixth-order PF algorithm with empirical bound, the TS algorithm, and the QSP algorithm (in both its segmented and non-segmented versions).

Observe that the TS algorithm outperforms the fourth-order PF algorithm with rigorous performance guarantees provided $n$ is at least about 30. However, the CNOT count over Clifford $+R_{z}$ is only negligibly improved, whereas the $T$ count over Clifford $+T$ is improved by a factor of about
7. This advantage comes at a significant cost in terms of space, as discussed in Section 1. For system sizes between 20 and 50 , the TS algorithm uses between 116 and 171 qubits, whereas the PF algorithm has no space overhead. Furthermore, the QSP algorithm always outperforms the TS algorithm with respect to both circuit size and number of qubits, so it is clearly preferred.

Recall that we implemented two variants of the QSP algorithm: the full algorithm (with random rotation angles in place of the true, hard-to-compute values) using the empirical Jacobi-Anger bound (193), and the segmented version using the rigorous error bound (190). Among all algorithms with a complete implementation and a rigorous accuracy guarantee, the segmented QSP algorithm has the lowest gate counts. With respect to the cnot count, it outperforms the fourth-order PF algorithm, and even the TS algorithm, by a factor of at least 5 . The improvement in the $T$ count is greater. The full QSP algorithm with empirical Jacobi-Anger bound has even lower gate complexity, but this will only be useful in practice if one can overcome the difficulty of classically computing the rotation angles. Furthermore, if we are willing to accept heuristic error estimates, then the PF algorithm significantly outperforms either QSP algorithm while also requiring no ancilla qubits.

## I. 3 Circuit optimization

To reduce the resource requirements for quantum simulation as much as possible, we post-processed all of our circuits with an automated circuit optimizer [57] as described in Section E.3. Figure 12 compares the gate counts of the PF, TS, and QSP algorithms before and after this optimization. The CNOT and $T$ counts of the fourth-order PF algorithm are both reduced by about $30 \%$ throughout the range of interest. While the gate counts of the TS and QSP algorithms are also reduced, the improvement is much less significant.

For the Clifford $+R_{z}$ resource estimates (see the left plot in Figure 12), the PF circuits show a $33 \%$ reduction in CNOT count, whereas both the TS and QSP circuits admit only about $1 \%$ reduction in their cnot counts. We expect to see improvements in the cnot count reduction for the TS and QSP circuits when using heavy instead of light optimization, although preliminary results suggest that this reduction will not qualitatively change the relative performance of the different algorithms. For the Clifford $+T$ resource estimates (see the right plot in Figure 12), the $T$ count reduction is approximately $30 \%$ for the PF algorithm, $0.5 \%$ for TS, and $1 \%$ for QSP.

With respect to CNOT count, whereas the TS algorithm outperforms the best PF algorithm with rigorous performance guarantees for all system sizes shown in Figure 12, optimization improves the PF algorithm to outperform the TS algorithm for system sizes smaller than about 30. While the TS algorithm is dominated by the more efficient QSP approach, this example nevertheless shows that optimization can sometimes affect the relative performance of algorithms.

We also compare the optimization of PF algorithms of different orders, as shown in Figure 13. We plot gate counts for PF circuits with the empirical bound, which offers the best performance. Since the structure of the PF circuit is not affected by choosing a different error bound, essentially the same relative improvements hold for other bounds as well.

The first-order PF circuits do not admit any optimization in the CnOt counts. The remaining PF circuits admit CNOT gate count reduction of about $33 \%$, with marginally more savings observed for higher orders (although this additional gain is too small to see in the scale used in Figure 13). The $R_{z}$ count of the PF circuits behaves similarly under optimization. For the first-order PF circuits, the $R_{z}$ count is not reduced. However, for all higher-order circuits, the optimizer can take advantage of the reflection symmetry in formulas of order 2 and higher, which reduces the $R_{z}$ counts in the original Clifford $+R_{z}$ circuits. The observed reduction is about $29 \%$.


Figure 12: Gate counts before (diamonds) and after (circles) circuit optimization. Left: CNOT gates over Clifford $+R_{z}$. Right: $T$ gates over Clifford $+T$.


Figure 13: Gate counts before (diamonds) and after (circles) circuit optimization for the first, second, fourth and sixth order PF algorithm with empirical bound. Left: CNOT gates over Clifford $+R_{z}$. Right: $R_{z}$ gates over Clifford $+R_{z}$.

## References

[1] Dorit Aharonov and Amnon Ta-Shma, Adiabatic quantum state generation and statistical zero knowledge, Proceedings of the 35th ACM Symposium on Theory of Computing, pp. 20-29, 2003, quant-ph/0301023.
[2] Matthew Amy, Dmitri Maslov, Michele Mosca, and Martin Roetteler, A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits, IEEE Transactions on ComputerAided Design of Integrated Circuits and Systems 32 (2013), no. 6, 818-830, arXiv:1206.0758.
[3] Ryan Babbush, Dominic W. Berry, Ian D. Kivlichan, Annie Y. Wei, Peter J. Love, and Alán Aspuru-Guzik, Exponentially more precise quantum simulation of fermions in second quantization, New Journal of Physics 18 (2016), no. 3, 033032, arXiv:1506.01020.
[4] Ryan Babbush, Dominic W. Berry, Yuval Sanders, Ian D. Kivlichan, Artur Scherer, Annie Wei, Peter Love, and Alán Aspuru-Guzik, Exponentially more precise quantum simulation of
fermions in the configuration interaction representation, to appear in Quantum Science and Technology, arXiv:1506.01029.
[5] Ryan Babbush, Jarrod McClean, Dave Wecker, Alán Aspuru-Guzik, and Nathan Wiebe, Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation, Physical Review A 91 (2015), 022311, arXiv:1410.8159.
[6] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A Megrant, E Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, E. Solano, and John M. Martinis, Digital quantum simulation of fermionic models with a superconducting circuit, Nature Communications 6 (2015), 7654, arXiv:1501.07703.
[7] Bela Bauer and Chetan Nayak, Analyzing many-body localization with a quantum computer, Physical Review X 4 (2014), no. 4, 041021, arXiv:1407.1840.
[8] Stephane Beauregard, Circuit for Shor's algorithm using $2 n+3$ qubits, Quantum Information and Computation 3 (2003), no. 2, 175-185, arXiv:quant-ph/0205095.
[9] David Beckman, Amalavoyal N. Chari, Srikrishna Devabhaktuni, and John Preskill, Efficient networks for quantum factoring, Physical Review A 54 (1996), 1034-1063, arXiv:quant-ph/9602016.
[10] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, arXiv:1707.04344.
[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 and Andrew M. Childs, Black-box Hamiltonian simulation and unitary implementation, Quantum Information and Computation 12 (2012), no. 1-2, 29-62, arXiv:0910.4157.
[13] 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 ACM Symposium on Theory of Computing, pp. 283-292, 2014, arXiv:1312.1414.
[14] 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.
[15] Dominic W. Berry, Andrew M. Childs, and Robin Kothari, Hamiltonian simulation with nearly optimal dependence on all parameters, Proceedings of the 56th IEEE Symposium on Foundations of Computer Science, pp. 792-809, 2015, arXiv:1501.01715.
[16] 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.
[17] Jonathan M. Borwein and Peter B. Borwein, Pi and the AGM, Wiley, 1987.
[18] Sergei Bravyi and Alexei Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, Physical Review A 71 (2005), 022316, arXiv:quant-ph/0403025.
[19] Kenneth R. Brown, Robert J. Clark, and Isaac L. Chuang, Limitations of quantum simulation examined by simulating a pairing Hamiltonian using nuclear magnetic resonance, Physical Review Letters 97 (2006), 050504, arXiv:quant-ph/0601021.
[20] Earl Campbell, Shorter gate sequences for quantum computing by mixing unitaries, Physical Review A 95 (2017), 042306.
[21] Yu Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J.Y. Mutus, P.J.J. O'Malley, C.M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T.C. White, Michael R. Geller, A.N. Cleland, and John M. Martinis, Qubit architecture with high coherence and fast tunable coupling, Physical Review Letters 113 (2014), 220502, arXiv:1402.7367.
[22] Donny Cheung, Dmitri Maslov, Jimson Mathew, and Dhiraj K. Pradhan, On the design and optimization of a quantum polynomial-time attack on elliptic curve cryptography, Quantum Information and Computation 9 (2009), no. 7-8, 610-621, arXiv:0710.1093.
[23] Andrew M. Childs, Quantum information processing in continuous time, Ph.D. thesis, Massachusetts Institute of Technology, 2004.
[24] Andrew M. Childs, On the relationship between continuous- and discrete-time quantum walk, Communications in Mathematical Physics 294 (2010), no. 2, 581-603, arXiv:0810.0312.
[25] Andrew M. Childs and Robin Kothari, Simulating sparse Hamiltonians with star decompositions, Proceedings of the 5th Conference on Theory of Quantum Computation, Communication, and Cryptography (TQC 2010), Lecture Notes in Computer Science, vol. 6519, pp. 94-103, 2011, arXiv:1003.3683.
[26] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su, https://github.com/njross/simcount.
[27] Andrew M. Childs and Nathan Wiebe, Hamiltonian simulation using linear combinations of unitary operations, Quantum Information and Computation 12 (2012), no. 11-12, 901-924, arXiv:1202.5822.
[28] Shantanu Debnath, Norbert M. Linke, Caroline Figgatt, Kevin A Landsman, Kevin Wright, and Christopher Monroe, Demonstration of a small programmable quantum computer with atomic qubits, Nature 536 (2016), 63-66, arXiv:1603.04512.
[29] Guillaume Duclos-Cianci and David Poulin, Reducing the quantum-computing overhead with complex gate distillation, Physical Review A 91 (2015), 042315, arXiv:1403.5280.
[30] Richard P. Feynman, Simulating physics with computers, International Journal of Theoretical Physics 21 (1982), no. 6-7, 467-488.
[31] Austin G. Fowler, Simon J. Devitt, and Lloyd C. L. Hollenberg, Implementation of Shor's algorithm on a linear nearest neighbour qubit array, Quantum Information and Computation 4 (2004), 237-251, arXiv:quant-ph/0402196.
[32] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron, Quipper: A scalable quantum programming language, ACM SIGPLAN Notices 48 (2013), no. 6, 333-342, arXiv:1304.3390.
[33] Thomas Häner, Martin Roetteler, and Krysta M. Svore, Factoring using $2 n+2$ qubits with Toffoli based modular multiplication, arXiv:1611.07995.
[34] Thomas Häner and Damian S. Steiger, 0.5 petabyte simulation of a 45-qubit quantum circuit, arXiv:1704.01127.
[35] Aram W. Harrow and Ashley Montanaro, Quantum computational supremacy, Nature 549 (2017), 203-209.
[36] Matthew B. Hastings, Dave Wecker, Bela Bauer, and Matthias Troyer, Improving quantum algorithms for quantum chemistry, Quantum Information and Computation 15 (2015), 1, arXiv:1403.1539.
[37] IARPA Quantum Computer Science program, https://www.iarpa.gov/index.php/research-programs/qcs.
[38] IBM, The Quantum Experience, http://www.research.ibm.com/ibm-q/, 2016.
[39] Stephen P. Jordan, Keith S. M. Lee, and John Preskill, Quantum algorithms for fermionic quantum field theories, arXiv:1404.7115.
[40] Stephen P. Jordan, Keith S. M. Lee, and John Preskill, Quantum algorithms for quantum field theories, Science 336 (2012), no. 6085, 1130-1133, arXiv:1111.3633.
[41] Stephen P. Jordan, Keith S. M. Lee, and John Preskill, Quantum computation of scattering in scalar quantum field theories, Quantum Information and Computation 14 (2014), 1014-1080, arXiv:1112.4833.
[42] Ivan Kassal, James D. Whitfield, Alejandro Perdomo-Ortiz, Man-Hong Yung, and Alán Aspuru-Guzik, Simulating chemistry using quantum computers, Annual Review of Physical Chemistry 62 (2011), no. 1, 185-207, arXiv:1007.2648.
[43] Jonas A. Kjäll, Jens H. Bardarson, and Frank Pollmann, Many-body localization in a disordered quantum Ising chain, Physical Review Letters 113 (2014), no. 10, 107204, arXiv:1403.1568.
[44] Thorsten Kleinjung, Kazumaro Aoki, Jens Franke, Arjen K. Lenstra, Emmanuel Thomé, Joppe W. Bos, Pierrick Gaudry, Alexander Kruppa, Peter L. Montgomery, Dag Arne Osvik, Herman te Riele, Andrey Timofeev, and Paul Zimmermann, Factorization of a 768-bit RSA modulus, Proceedings of the 30th Annual Cryptology Conference (CRYPTO), pp. 333-350, 2010.
[45] Vadym Kliuchnikov, Dmitri Maslov, and Michele Mosca, Fast and efficient exact synthesis of single-qubit unitaries generated by Clifford and $T$ gates, Quantum Information and Computation 13 (2013), no. 7-8, 607-630.
[46] Noboru Kunihiro, Exact analyses of computational time for factoring in quantum computers, IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences 88 (2005), no. 1, 105-111.
[47] Samuel A. Kutin, Shor's algorithm on a nearest-neighbor machine, arXiv:quant-ph/0609001.
[48] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos, Universal digital quantum simulation with trapped ions, Science 334 (2011), no. 6052, 57-61, arXiv:1109.1512.
[49] Norbert M. Linke, Dmitri Maslov, Martin Roetteler, Shantanu Debnath, Caroline Figgatt, Kevin A. Landsman, Kenneth Wright, and Christopher Monroe, Experimental comparison of two quantum computing architectures, Proceedings of the National Academy of Sciences 114 (2017), no. 13, 3305-3310, arXiv:1702.01852.
[50] Seth Lloyd, Universal quantum simulators, Science 273 (1996), no. 5278, 1073-1078.
[51] Guang Hao Low and Isaac L. Chuang, Hamiltonian simulation by qubitization, arXiv:1610.06546.
[52] Guang Hao Low and Isaac L. Chuang, Optimal Hamiltonian simulation by quantum signal processing, Physical Review Letters 118 (2017), 010501, arXiv:1606.02685.
[53] Guang Hao Low, Theodore J. Yoder, and Isaac L. Chuang, Methodology of resonant equiangular composite quantum gates, Physical Review X 6 (2016), 041067, arXiv:1603.03996.
[54] David J. Luitz, Nicolas Laflorencie, and Fabien Alet, Many-body localization edge in the random-field Heisenberg chain, Physical Review B 91 (2015), no. 8, 081103, arXiv:1411.0660.
[55] Dmitri Maslov, Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization, Physical Review A 93 (2016), 022311, arXiv:1508.03273.
[56] Dmitri Maslov, Optimal and asymptotically optimal NCT reversible circuits by the gate types, Quantum Information and Computation 16 (2016), no. 13-14, 1096-1112, arXiv:1602.02627.
[57] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov, Automated optimization of large-scale quantum circuits with continuous parameters, 2017, arXiv:1710.07345.
[58] Rahul Nandkishore and David A. Huse, Many-body localization and thermalization in quantum statistical mechanics, Annual Review of Condensed Matter Physics 6 (2015), no. 1, 15-38, arXiv:1404.0686.
[59] Arijeet Pal and David A. Huse, Many-body localization phase transition, Physical Review B 82 (2010), no. 17, 174411, arXiv:1010.1992.
[60] Archimedes Pavlidis and Dimitris Gizopoulos, Fast quantum modular exponentiation architecture for Shor's factorization algorithm, Quantum Information and Computation 14 (2014), no. 7-8, 649-682, arXiv:1207.0511.
[61] Edwin Pednault, John A. Gunnels, Giacomo Nannicini, Lior Horesh, Thomas Magerlein, Edgar Solomonik, and Robert Wisnieff, Breaking the 49-qubit barrier in the simulation of quantum circuits, arXiv:1710.05867.
[62] 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.
[63] Sadegh Raeisi, Nathan Wiebe, and Barry C. Sanders, Quantum-circuit design for efficient simulations of many-body quantum dynamics, New Journal of Physics 14 (2012), 103017, arXiv:1108.4318.
[64] Markus Reiher, Nathan Wiebe, Krysta M. Svore, Dave Wecker, and Matthias Troyer, Elucidating reaction mechanisms on quantum computers, Proceedings of the National Academy of Sciences 114 (2017), no. 29, 7555-7560, arXiv:1605.03590.
[65] Martin Roetteler, Michael Naehrig, Krysta M. Svore, and Kristin Lauter, Quantum resource estimates for computing elliptic curve discrete logarithms, arXiv:1706.06752.
[66] Neil J. Ross and Peter Selinger, Optimal ancilla-free Clifford $+T$ approximation of $z$-rotations, Quantum Information and Computation 16 (2016), no. 11-12, 901-953, arXiv:1403.2975.
[67] Michael Schreiber, Sean S. Hodgman, Pranjal Bordia, Henrik P. Lüschen, Mark H. Fischer, Ronen Vosk, Ehud Altman, Ulrich Schneider, and Immanuel Bloch, Observation of manybody localization of interacting fermions in a quasirandom optical lattice, Science 349 (2015), no. 6250, 842-845, arXiv:1501.05661.
[68] M. Serbyn, M. Knap, S. Gopalakrishnan, Z. Papić, N. Y. Yao, C. R. Laumann, D. A. Abanin, M. D. Lukin, and E. A. Demler, Interferometric probes of many-body localization, Physical Review Letters 113 (2014), no. 14, 147204, arXiv:1403.0693.
[69] Vivek V. Shende, Stephen S. Bullock, and Igor 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.
[70] Jacob Smith, Aaron Lee, Philip Richerme, Brian Neyenhuis, Paul W. Hess, Philipp Hauke, Markus Heyl, David A. Huse, and Christopher Monroe, Many-body localization in a quantum simulator with programmable random disorder, Nature Physics 12 (2016), 907-911, arXiv:1508.07026.
[71] Rolando Somma, Gerardo Ortiz, James E. Gubernatis, Emanuel Knill, and Raymond Laflamme, Quantum algorithms for fermionic simulations, Physical Review A 64 (2002), 022319, arXiv:cond-mat/0012334.
[72] Rolando Somma, Gerardo Ortiz, James E. Gubernatis, Emanuel Knill, and Raymond Laflamme, Simulating physical phenomena by quantum networks, Physical Review A 65 (2002), 042323, arXiv:quant-ph/0108146.
[73] Chao Song, Kai Xu, Wuxin Liu, Chui ping Yang, Shi-Biao Zheng, Hui Deng, Qiwei Xie, Keqiang Huang, Qiujiang Guo, Libo Zhang, Pengfei Zhang, Da Xu, Dongning Zheng, Xiaobo Zhu, H. Wang, Y.-A. Chen, C.-Y. Lu, Siyuan Han, and Jian-Wei Pan, 10-qubit entanglement and parallel logic operations with a superconducting circuit, Physical Review Letters 119 (2017), 180511, arXiv:1703.10302.
[74] Masuo Suzuki, General theory of fractal path integrals with applications to many-body theories and statistical physics, Journal of Mathematical Physics 32 (1991), no. 2, 400-407.
[75] Yasuhiro Takahashi and Noboru Kunihiro, A quantum circuit for Shor's factoring algorithm using $2 n+2$ qubits, Quantum Information and Computation 6 (2006), no. 2, 184-192.
[76] Dave Wecker, Bela Bauer, Bryan K. Clark, Matthew B. Hastings, and Matthias Troyer, Gate count estimates for performing quantum chemistry on small quantum computers, Physical Review A 90 (2014), 022305, arXiv:1312.1695.
[77] Dave Wecker, Matthew B. Hastings, Nathan Wiebe, Bryan K. Clark, Chetan Nayak, and Matthias Troyer, Solving strongly correlated electron models on a quantum computer, Physical Review A 92 (2015), 062318, arXiv:1506.05135.
[78] Jiehang Zhang, Guido Pagano, Paul W. Hess, Antonis Kyprianidis, Patrick Becker, Harvey Kaplan, Alexey V. Gorshkov, Zhe-Xuan Gong, and Christopher Monroe, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, arXiv:1708.01044.


[^0]:    ${ }^{1}$ Recent work has demonstrated simulation of 56 -qubit computations, but only for circuits of much smaller depth than those considered in our work [61].

[^1]:    ${ }^{2}$ Reference [47] does not give explicit resource counts; we estimate them as described in Appendix A.

[^2]:    ${ }^{3}$ While a construction with this property was not explicitly developed in [55], it appears that any multiplycontrolled Toffoli gate with arbitrary positive and negative controls can be implemented with the same number of gates as the multiply-controlled Toffoli gate of the same size with only positive controls.

