ARTICLE OPEN

Low-cost quantum circuits for classically intractable instances of the Hamiltonian dynamics simulation problem

Yunseong Nam¹ and Dmitri Maslov²

We develop circuit implementations for digital-level quantum Hamiltonian dynamics simulation algorithms suitable for implementation on a reconfigurable quantum computer, such as trapped ions. Our focus is on the codesign of a problem, its solution, and quantum hardware capable of executing the solution at the minimal cost expressed in terms of the quantum computing resources used, while demonstrating the solution of an instance of a scientifically interesting problem that is intractable classically. The choice for Hamiltonian dynamics simulation is due to the combination of its usefulness in the study of equilibrium in closed quantum mechanical systems, a low cost in the implementation by quantum algorithms, and the difficulty of classical simulation. By targeting a specific type of quantum computer and tailoring the problem instance and solution to suit physical constraints imposed by the hardware, we are able to reduce the resource counts by a factor of 10 in a physical-level implementation and a factor of 30–60 in a fault-tolerant implementation over state-of-the-art.

npj Quantum Information (2019)5:44; <https://doi.org/10.1038/s41534-019-0152-0>

INTRODUCTION

Quantum supremacy is a computational experiment designed to demonstrate a computational capability of a quantum machine that cannot be matched by a classical computer. It is highly relevant to this paper, since we too focus on a quantum computation of the size that cannot be performed via classical means. Quantum supremacy is an important milestone in the development of quantum computers. Multiple IT giants are targeting quantum supremacy (IBM: <https://spectrum.ieee.org/tech-talk/computing/hardware/ibm-edges-closer-to-quantum-supremacy-with-50qubit-processor> (posted 2017-11-15), Intel: <https://spectrum.ieee.org/tech-talk/computing/hardware/intels-49qubit-chip-aims-for-quantum-supremacy> (posted 2018-01-09), and Google: <https://spectrum.ieee.org/computing/hardware/google-plans-to-demonstrate-the-supremacy-of-quantum-computing> (posted 2017-05-24) race towards achieving quantum supremacy), and it is perhaps reasonable to anticipate that a successful demonstration may be obtained within at most a few years.

Once quantum supremacy is demonstrated, a next step is to go beyond what the supremacy would have achieved. The supremacy experiment proposed in ref. ¹, in particular, reduces to the execution of a random quantum circuit on a quantum computer that is too large for a classical computer to cope with simulating. In this work, rather than focusing on an artificial problem designed purely for demonstrating quantum supremacy,¹ we target the selection of a known computational problem and a specific input instance, such that, to the best of our knowledge, the problem/instance pair requires a classically intractable computation. We develop an optimized quantum circuit computing the answer for the selected problem/instance pair that can be suitable for the execution on near-term quantum computers. Our overarching goal is to select a problem/instance pair and develop a short enough circuit such that it will be the smallest among all circuits

solving a post-classical instance of a scientific problem. A quantum computation described by such circuit constitutes a qualitative step forward, where a quantum computer can now be thought of as being a tool in the solution of a problem rather than the focus of the study. A more advanced demonstration past the one we are reporting in this work could target the solution of a problem with a commercial value.

In the “Results” section, we report the quantum resources required for running the proposed post-supremacy experiments and discuss the quality of our results, expressed in terms of quantum circuit gate count and depth, and viewed through the lens of comparisons to prior and similar-spirited work. We summarize and discuss our overall findings in the “Discussion” section. In the “Methods” section, we introduce the details of the problem and the specific instance of this problem that we are proposing as satisfying the conditions outlined in the previous paragraph. We prove that the type of problem we consider is classically hard under the assumption that Polynomial Hierarchy does not collapse. We also describe our solution, including numerous techniques used to improve quantum computing resources used. We stress that we develop complete and fully specified quantum circuits as a part of this study.

RESULTS

Table 1 reports gate counts in the post-supremacy Hamiltonian dynamics simulation with Heisenberg interactions and random disorders on graphs with small diameter (readers are strongly encouraged to read the section “Methods” for technical description of the Hamiltonian and the optimized implementation detail of the simulation, used throughout this section). We show resource counts targeted for both physical-level and fault-tolerant implementations. We note that the counts were obtained while optimizing the circuit depth. We found that they may be

¹IonQ, College Park, MD 20740, USA and ²National Science Foundation, Alexandria, VA 22314, USA
Correspondence: Yunseong Nam (nam@ionq.co) or Dmitri Maslov (dmitri.maslov@gmail.com)

Received: 7 August 2018 Accepted: 1 April 2019

Published online: 27 May 2019

Table 1. Gate counts for our proposed experiments

Graph	Pre-fault tolerant				Fault tolerant (Gridsynth)			Fault tolerant (RUS)		
	n	CNOT	R_z	Depth	n	CNOT	τ	n	CNOT	τ
4th order										
(3–5–70)	70	648,885	700,367	25,333	126	4,751,775	13,398,840	131	9,573,795	6,751,395
(4–4–98)	98	1,590,750	1,696,926	41,811	179	11,041,734	23,878,140	185	19,079,280	12,615,702
(5–3–72)	72	1,924,560	2,007,822	66,525	142	13,568,294	28,047,884	148	22,865,870	15,011,834
6th order										
(3–5–70)	70	1,873,350	2,021,797	72,587	126	9,802,780	26,068,380	131	18,995,860	13,210,730
(4–4–98)	98	6,011,250	6,412,126	155,259	179	37,283,365	77,435,190	185	63,102,200	41,364,575
(5–3–72)	72	4,519,800	4,715,202	154,005	142	24,882,660	49,081,320	148	40,935,060	26,578,440
(7–2–50)	50	1,124,700	1,160,419	51,795	96	6,991,712	12,887,472	101	11,154,260	7,172,048

Fig. 1 Quantum resource counts as a function of the number of logical qubits for various quantum circuits. Solid plot symbols denote physical-level/pre-fault tolerant (pFT) circuits. Hollow plot symbols denote fault-tolerant (FT) circuits. For pFT circuits, we use CNOT gate counts as the resource requirement and the logical qubit count is the same as the bare physical qubit count. For FT circuits, we use τ gate counts as the resource requirement. For the various similarly spirited previous work, see³ for Shor,² for FeMoco,⁴ for Hubbard and Jellium, and⁵ for Heisenberg Hamiltonian over cycle (C). The 4th order formula results reported in this paper (Table 1) appear as “Heisenberg” in the figure

improved by about 20% if we optimize the gate counts themselves instead.

For the purpose of comparison to prior, similarly spirited work reporting detailed gate counts (see Fig. 1 for a visual representation of a more comprehensive list of data), we focus on our best result, the Hamiltonian simulation over the (3–5–70) graph (see the section “Methods” for detail). In quantum chemistry, the simulation of FeMoco (the primary cofactor of nitrogenase, which is an enzyme used in the nitrogen fixation) required the circuit with 10^{14} τ gates over 111 qubits.² In comparison, our fault-tolerant circuits simulating a Heisenberg Hamiltonian system with 6.8×10^6 and 1.3×10^7 τ gates (4th and 6th order product formulas, correspondingly) are significantly shorter, while relying on a comparable number of qubits, 131, and corresponding to solving a task of a similar classical complexity. To factor a 1024 digit integer number,³ constructed a circuit with 5.7×10^9 τ gates spanning 3132 qubits. Our circuits simulating a Heisenberg Hamiltonian system are orders of magnitude shorter and operate over a much smaller number of qubits. Recent work⁴ required 10^8 τ gates to solve a problem in the study of solid-state electronic structure, whereas our τ count is only 6.8×10^6 . Finally, the task of a very similar complexity (70 qubits, $\epsilon = 0.001$, Hamiltonian on a cycle⁵) is solved using 4.2×10^8 τ gates or 6.7×10^6 CNOT gates,

showing the advantage of our approach by a factor of 60 in the τ count and a factor of 10 in the CNOT count. In fact, the circuits we developed to simulate the Hamiltonian evolution are so short that we hope they may be possible to execute on pre-fault tolerant quantum computers.

The two-qubit depth of our quantum circuits is very small, making them particularly suitable for implementations over quantum-information processors (QIPs) with limited T_1/T_2 coherence times. Our shortest circuit has the two-qubit gate depth of only 25,333 (Table 1), well within the limits of best trapped ions coherence time of 10 min⁶ divided by the two-qubit gate operating time of about 250 μ s.⁷ The fidelity of physical-level two-qubit gates may not be required to be as low as a simple formula, $\frac{1}{\sqrt{\#gates}}$, prescribes, before a computation with this many gates can be executed. Indeed, in cases when the dominant error source is of the random under-/over-rotation kind, the error in the execution of the entire circuit may scale sublinearly as a function of the number of gates, and potentially as low as the square root of the number of gates (this is in contrast to additive errors, such as those resulting from the decoherence), and thus we can expect to be able to execute a larger number of gates with a given two-qubit gate fidelity. Sublinear (in the number of gates) scaling of the accumulated error in a quantum computation has already been witnessed experimentally,⁷ despite small circuit sizes considered. Applying the mixing unitaries approach⁸ to the physical-level circuits through drawing a pulse sequence from a probability distribution given by multiple approximations that mix to a high-quality two-qubit gate can further lower the requirements for the two-qubit gate error necessary to carry out the computation we proposed. The lower bound on the gate fidelity of $\frac{1}{\sqrt{\#gates}}$, being about 0.001 for $\#gates = 1,000,000$ was already achieved experimentally,⁹ suggesting that the two-qubit gate fidelity required to carry out the proposed computations may be within reach. Finally, the number of qubits available in either superconducting circuits (72) or trapped-ion (79) technology readily meets the qubit count requirement in our circuits (Google announced a 72-qubit machine (posted 2018-3-5), IonQ announced a 79-qubit machine (posted 2018-12-11)).

We believe circuit implementations reported in this paper could be particularly relevant to the near-term QIPs based on the trapped-ion technology. This is because the low-diameter graphs, suitable for the kinds of experiments we considered, have connectivities that require long-range interactions between qubits. For 50 or larger number of qubits, a trapped-ion quantum simulator is already known to be capable of all-to-all coupling^{10,11} and there already exist a suite of systematic methods to generate a control signal that implements an arbitrarily connected two-qubit gate (see, for instance,¹²). Leveraging all-to-all connectivity of the trapped ions QIP, we expect no overhead cost in shuttling quantum information around, a serious point to consider, as has been pointed out in ref.⁷

DISCUSSION

In this paper, we synthesized short quantum circuits that aim to solve a scientifically interesting problem. Specifically, we considered the Heisenberg Hamiltonian simulation with a random disorder on a suite of graphs with small diameter. Compared with the previous state-of-the-art, our work shows significant gate count savings, a short circuit depth, while natively relying on the qubit-to-qubit connectivity suitable for the implementation over a reconfigurable quantum computer, such as the trapped ions one. Specifically, we reported a circuit for simulating Hamiltonian dynamics over (3–5–70) graph for the time $t = 10$ and accurate to within the error $\varepsilon = 0.001$ using at most 648,885 CNOT gates in depth 25,333 in a pre-fault tolerant implementation and 6,751,395 τ gates in a fault-tolerant implementation. We believe the problem instance we considered is intractable for classical computers and yet the quantum resource estimates are very low, showing the promise for solving interesting problems by a quantum computer in a not-too-distant future.

METHODS

Problem

We consider the problem of simulating Hamiltonian dynamics for time t , accurate to within the error ε , where the target Hamiltonian H is defined as follows:

$$H := \sum_{(i,j) \in E(G)} (\vec{\sigma}_x^i \vec{\sigma}_x^j + \vec{\sigma}_y^i \vec{\sigma}_y^j + \vec{\sigma}_z^i \vec{\sigma}_z^j) + \sum_{i=1}^n d_i \vec{\sigma}_z^i, \quad (1)$$

where $\vec{\sigma}_x^i, \vec{\sigma}_y^i$, and $\vec{\sigma}_z^i$ denote Pauli x, y , and z matrices acting on the qubit i , $d_i \in [-1, 1]$ are chosen uniformly at random, G is a graph describing the two-qubit interactions, $E(G)$ is the set of its edges, and n is the number of qubits. Such Hamiltonian is known as the Heisenberg Hamiltonian over a graph G with a random disorder in the Z direction. It has been studied in refs. ^{13–15} in the context of many-body localization.

We chose n to be in the range 50–100. This is because the largest quantum circuit (vector state as opposed to full unitary) simulations demonstrated to date were achieved with dozens of qubits over a circuit depth of about 40.^{16,17} This means that with anywhere more than 50 qubits and depth in excess of, say 200, the problem of quantum circuit simulation may likely become intractable for a classical computer with the simulation techniques such as refs. ^{16,17} Note that our circuits, despite numerous optimizations applied, remain substantively deeper than those considered in refs.,^{16,17} motivating our choice to consider reducing the number of qubits from as many as 144 for a shallow circuit with depth 27¹⁶ down to 50 at the cost of significantly extending the anticipated length of the computation. We also note that the underlying qubit-to-qubit connectivity pattern in our circuits does not appear to allow breaking the qubit interaction graph into two components by cutting a small number of edges, which lies at the core of efficient simulations such as refs. ^{16,17}. The smallest number of qubits we chose to consider, 50, exceeds the number 42 used in the best high-depth state-vector type simulation.¹⁸ Finally, we highlight that the largest numerical simulation of the kind of Hamiltonian we consider is restricted to just 22 qubits,¹³ whereas our work focuses on the Hamiltonians over at least 50 qubits.

We chose the evolution time $t = 2d$, where d is the diameter of graph G , similarly to ref. ⁵ The motivation behind such choice is as follows: it takes time at most d ¹⁹ for quantum information to propagate from any node in the underlying graph G to any other node, as such, one may expect to enter a highly entangled simulation regime by the simulation time of d . This means that our simulation spends at least half the time evolving in the regime that we believe is difficult to simulate classically. Note that due to the use of the product formula approach^{20,21} in our simulations, circuits for other selection of time t can be developed as effortlessly as changing the number of times a certain block operation is applied.

A previous study⁵ showed that the product formula algorithm(s)^{20,21} for Hamiltonian simulation with a heuristic bound yields best practical results. We note that choosing a small diameter graph G , such as what we do next, is natural, given the discussion in the previous paragraph. However, small diameter graphs require large number of edges, and the number of edges in the graph G is directly proportional to the circuit complexity of the single stage of product formula. Thus, the balance between graph diameter and the number of edges has to be chosen carefully so as to

minimize quantum computational resources while maximizing the expected classical difficulty of the simulation.

We chose graph G to be a minimal distance k -regular graph over n nodes. We select k such that the effort required to develop the individual two-qubit gates for all $\frac{nk}{2}$ interactions prescribed by the graph in a technology such as the trapped ions is not large,²² while keeping the graph distance small (resulting in a short enough time t for the evolution before Hamiltonian dynamics simulation enters a regime that is believed to be classically difficult) and the number of qubits as large as possible (while keeping it to between 50 and 100). In practice, we selected the following values of k : 3, 4, 5, 6, and 7. The specific regular graphs G considered in our work can be described by the respective degree-diameter-nodes triple $(k-d-n)$, as follows: (3–5–70) Alegre-Fiol-Yebra graph,²³ (4–4–98) graph by Exoo,²⁴ (5–3–72) graph by Exoo,²⁴ and (7–2–50) Hoffman–Singleton graph.²⁵ The largest known 6-regular distance-2 graph has 32 nodes (less than 50 targeted in our work), and the largest known 6-regular distance-3 graph has 110 nodes (more than 100 targeted in this work).²⁶ Therefore, we did not consider 6-regular graphs. Note that our goal was to select a graph with a large number of nodes, small degree, and small distance. A plenty of such graphs can be developed so long as one chooses degree and distance parameters above the minimums known for a given n , selects n smaller than the known maximum for the fixed degree and diameter, and expands the attention to graphs other than the regular kind. Our selection of graphs is very restrictive so as to narrow down the set of specific graphs explicitly considered in our work to a few. We believe the performance of the Hamiltonian simulation over similar graphs to those considered can be similar.

We chose the approximation error $\varepsilon = 0.001$, measured as the spectral norm distance between the target ideal evolution and the evolution obtained by running our quantum simulation circuits, as in ref. ⁵ This choice of the error value is largely arbitrary, but some choice is necessary to explicitly construct the circuits.

Classical hardness

We next show that the type of Hamiltonian dynamics simulation problem considered in our work is difficult for classical computers under the assumption that Polynomial Hierarchy does not collapse.²⁷ Thus, it is likely that quantum computers will have to be used to solve sufficiently large instances of the respective problem (e.g., sampling from the distribution given by the evolution of the given Hamiltonian).

To prove classical hardness of dynamics simulation of the type of Hamiltonian studied in our work (1), consider this Hamiltonian to be an instance drawn from a larger set, $\mathcal{H}$, being the set of Hamiltonians with time-dependent Heisenberg interactions that can be controllably turned on and off, and a random disorder that can apply in any fixed direction for a given qubit, X, Y , or Z , and can also be turned on and off controllably with time. Each Hamiltonian in the set $\mathcal{H}$ applies qubit-to-qubit interactions as determined by the respective small-degree small-diameter graph it is considered over.

The specific instance of the Hamiltonian dynamics simulation problem studied in this work constitutes the hardest instance to simulate in the set $\mathcal{H}$ using a quantum digital computer with the algorithms considered. This is because all two-body interactions are always on and thus, due to the use of Suzuki–Trotter approach, require a maximal number of gates to be simulated; note that when an interaction is turned off, it takes no gates to simulate the evolution for such periods of time, thereby reducing the gate counts. Similarly, single-qubit disorders take no gates to simulate for those times when they are turned off (thereby resulting in the τ count reduction in the fault-tolerant case). Random disorders pointing in a given fixed direction X, Y , or Z does not change how various optimizations reduce the relevant gate counts; this is because commuting rotations over a fixed axis always add into a combined rotation, being the property that we base the resource reductions on. Furthermore, basis changes $R_x \mapsto R_z$ and $R_y \mapsto R_z$ used in the circuit decompositions are accomplished via the use of Hadamard and Phase gates, and do not affect any of the costing metrics considered.

To establish classical hardness,²⁷ we show that the Hamiltonians in the set $\mathcal{H}$ can be used to obtain universal evolutions. Specifically, we construct a computationally universal library consisting of arbitrary single-qubit and the CNOT gates. To assist with showing computational universality, consider a Hamiltonian such that all disorders but one point in the Z direction, and there is precisely one qubit q with disorder pointing in the Y direction.

Fig. 2 Circuit to construct a CNOT gate

Fig. 3 CNOT count optimal implementation of the Heisenberg interaction up to a global phase of $e^{-i\pi/4}$. This is a special case of the circuit shown in Fig. 6 of ref. ³⁶ Note that this circuit requires 3 CNOT gates

- **Arbitrary single-qubit gate:** To construct arbitrary single-qubit gate rely on Euler's angle decomposition in the form $R_y(\alpha)R_z(\beta)R_y(\gamma)$. First, consider applying arbitrary single-qubit rotation to the qubit q . R_y gate can be induced by directly evolving the disorder while all other disorders/interactions are turned off. To apply R_z , use Heisenberg interaction with the maximal rotation angle to SWAP qubit q into any of the neighboring qubits, apply R_z using the disorder, and finally SWAP back. To apply an arbitrary single-qubit gate to any other qubit, first SWAP it into qubit q (note that the length of the SWAP chain is at most logarithmic due to the choice of the underlying connectivity graph), and apply the above algorithm.
- **CNOT:** To construct the CNOT gate, use the following circuit (Fig. 2), relying on the operations implementable with the Hamiltonian considered.

Note that we proved classical hardness for a Hamiltonian that is slightly different from the one explicitly studied in this work. Our motivation was to select a Hamiltonian already used in the context of many-body localization,^{13–15} being a potential application of our study. We stress that the Hamiltonian used to prove classical intractability can be implemented with at most as many gates as the one we studied in our paper in the pre-fault tolerant case, and using fewer gates in the fault-tolerant case. This is because we can equalize the strengths of the single-qubit disorders and thus implement those with the input-weight algorithm discussed later in the paper, leading to additional reductions in the τ count.

Solution

Algorithm. Suzuki–Trotter formula-based algorithms with empirical bound showed themselves as a potent candidate among quantum algorithms that simulate Hamiltonian dynamics.^{5,20,21} We therefore chose to focus exclusively on this type of algorithms. Specifically, for a Hamiltonian of the form $H = \sum_j a_j H_j$, we approximate the evolution operator according to

$$\exp\left(-it \sum_{j=1}^L a_j H_j\right) \approx [S_{2k}(\lambda)]^r, \quad (2)$$

where $\lambda := -it/r$ and

$$S_2(\lambda) := \prod_{j=1}^L \exp(a_j H_j \lambda / 2) \prod_{j=L}^1 \exp(a_j H_j \lambda / 2),$$

$$S_{2k}(\lambda) := [S_{2k-2}(p_k \lambda)]^2 S_{2k-2}((1 - 4p_k)\lambda) [S_{2k-2}(p_k \lambda)]^2, \quad (3)$$

with $p_k := 1/(4 - 4^{1/(2k-1)})$ for $k > 1$.²⁰

We selected 4th ($k=2$) and 6th ($k=3$) order formula versions since in our case they achieved the best results. For a given order formula, the algorithm applies the operation S_4/S_6 r times, where S_4/S_6 contains 10/50 repetitions of the product of exponentials of the individual terms of the Hamiltonian H_i and r is the number of repetitions of S_4/S_6 for a given selection of the evolution time t , desired accuracy ε , underlying graph G , and the selection of random disorders d_i (1). We note that for a fixed ε , graph G , and a given selection of random disorders, r is proportional to a growing function of t . Explicit bounds on the value r are $O(t^{1+1/4})$ for the 4th order formula and $O(t^{1+1/6})$ for the 6th order formula.²¹

Codesign principle. Since the selection of t in our implementation is $t = 2d$, it is important to minimize the diameter of the underlying graph G , which explains our focus on the low-diameter graphs. The cost of S_4/S_6 can

Fig. 4 Implementation of the Heisenberg interaction, optimal in the number of real-valued degrees of freedom, up to a global phase of $e^{-i\pi/4}$. Implementation of the controlled- Z^a can be substituted from Fig. 5

Fig. 5 Ancilla-aided, measurement/feedforward-based fault-tolerant controlled- Z^a gate, imported from ref. ³⁷

be described as $10X/50X$, where X is the gate cost of the implementation of a single stage of the product of individual terms in the target Hamiltonian. In our work, we obtained physical-level implementations of the X stage with the cost of $\frac{3nk}{2}$ CNOT gates and the two-qubit depth $3k$ to $3k+3$, and fault-tolerant cost of $O(nk)$ τ gates and $O(nk)$ CNOT gates with depth $O(k(\log(n) + a))$, where parameter a is defined as the depth of the approximation of an R_z (see below for details). This means that all figures of merit are linearly dependent on the degree k of the graph G , and therefore to achieve the best performance, the degree k must be minimized.

Circuit design. To optimize the depth of our circuits, we employed a customized version of the implementation of Vizing's theorem,²⁸ which guarantees circuit depth k or $k+1$ for the implementation of a single stage of the product of exponentials of the target Hamiltonian. Our modification includes additional heuristic that randomly reorders the list of edges of the graph G in an attempt to find a depth- k layout when a depth $k+1$ layout is found, though in practice we found this to be of limited use. Our modification also readily accepts a manual input in case a depth- k layout is known, although we chose not to employ this method for the circuits we consider in this paper, since in general the problem of finding a depth- k layout, if at all exists, is NP-complete²⁹ and thus it is unlikely that the appropriate manual input would be known for a generic graph.

We implemented the Heisenberg interaction, $\exp[-ia(\vec{\sigma}_x^i \vec{\sigma}_x^j + \vec{\sigma}_y^i \vec{\sigma}_y^j + \vec{\sigma}_z^i \vec{\sigma}_z^j)]$, using two different circuits, depending on whether we focus on saving quantum resources in a pre-fault tolerant (physical-level) or a fault-tolerant implementation. In particular, for the physical-level implementation, we used the circuit shown in Fig. 3 in order to minimize the most expensive two-qubit gates, whereas for the fault-tolerant implementation, we used the circuit shown in Fig. 4 in order to minimize the most expensive τ gates. By directly synthesizing the Heisenberg interaction,

compared with the standard Pauli-matrix basis approach, for instance employed in ref. ⁵, we save 50% of the cost in the pre-fault tolerant implementation (evidenced through the reduction of the CNOT gate count from 6 down to 3) and almost 66% of the cost in the fault-tolerant implementation (evidenced through the reduction from 3 R_z gates down to 1 R_z and 4 τ gates). Because our construction directly implements the Heisenberg interaction, as opposed to an approximate implementation with XX, YY, and ZZ interactions that arise from the standard approach, our implementation also performs better at the algorithmic level, i.e., we do not need as large value of r as in the standard approach to keep the overall error level down that incurs from the approximate Pauli-basis implementation.

We laid out the Heisenberg interaction terms in the circuit implementation of the product formula algorithm in alternate orders—forward and reverse—to ensure we obtain maximal gain from the application of the circuit optimizer.³⁰ We furthermore modified the original optimizer³⁰ to ensure it can natively handle controlled- $z^a = cz^a$ gates and apply the merging rule $cz^a(x, y)cz^b(x, y) \mapsto cz^{a+b}(x, y)$. Also implemented was the merging of two Heisenberg interactions as per circuit implementation in Fig. 3 as an optimization rule, and the capability of being able to handle classically controlled quantum gates. The quality of optimization by the automated optimizer ranged from 7 to 14% in the CNOT gate count reduction, and 16 to 24% in the R_z count reduction, with the simulations over lower degree graphs yielding a better quality of optimization. Indeed, circuit implementations over graphs with lower degree have a larger proportion of gates at the edges of the circuit implementation of the Hamiltonian terms, and these are the types of gates that often admit optimizations by techniques such as ref. ³⁰

Fault-tolerant circuits. For fault-tolerant implementations, we decided to break the error budget evenly between the algorithmic errors that arise from the product formula algorithm and the approximation errors that arise from the approximation of R_z gates in the Clifford + τ basis. We distributed the approximation error budget evenly across all R_z gates in the given circuit. We employed three optimization strategies for fault-tolerant implementations: the mixing unitaries approach detailed in ref. ⁸, the application of equal-angle R_z rotations through computing input weight,³¹ and a combination of gridsynth^{32,33} and repeat-until-success (RUS)³⁴ strategies for the approximation of R_z gates by Clifford + τ circuits. These three strategies are detailed in the next four paragraphs.

The mixing unitaries approach⁸ relies on generating a set of four approximating circuits for a given $R_z(\theta)$ gate, and then applies a randomly drawn approximation from the set of four for every occurrence of the $R_z(\theta)$ gate in the circuit, subject to a certain probability distribution. We found out through simulation experiments that in practice the mixing unitaries approach achieves less than quadratic⁸ improvement in the error. This is because we measure the expected distance between the approximating circuit and the desired evolution, whereas⁸ measures the error $||\sum_k p_k U_k - V||$ between a collection of approximating circuits U_k given by the probability distribution $\{p_k\}$ and the target evolution V . Clearly, our metric is more restrictive, and in case when the probability distribution by a collection of approximating unitaries is an acceptable metric, further gate count reductions to those we reported in this paper become possible.

To study practical performance of the mixing strategy, we chose to investigate an amenable sample case of (5–1–6) graph in depth. We varied the per-gate error budget by tweaking the power p of $(\epsilon_{\text{approx}}/(N_{R_z}))^p$, where ϵ_{approx} is the approximation error budget and N_{R_z} is the number of R_z gates in the given circuit, since the number of τ gates for an approximation sequence scales linearly in the logarithm of the inverse of per-gate error level. An extensive numerical investigation showed that choosing p between 0.86 and 0.89 works well for the direct synthesis method and the values of p between 0.65 and 0.68 represent the advantage by the mixing approach. We therefore chose to use $p = (0.65 + 0.68)/2 = 0.665$ as the per-gate approximation power for the mixing unitaries method, for all cases we considered. We believe this is a safe assumption, since in the larger circuits the mixing strategy is expected to give better results as it takes more time to properly average out the errors. Employing the mixing unitaries strategy allows an estimated τ count savings of 25 to 33%, depending on whether the power $p = 0.875$ or $p = 1$ is considered as the starting point.

In our implementation of the two-qubit Hamiltonian interaction, we lay out the respective circuitry in parallel with the help of Vizing's theorem. This results in the parallel application of as many as $m \leq \frac{n}{2} R_z(\theta)$ gates with equal rotation angles [Note that the same idea can be applied to R_z gates representing random disorders and looking for rotations that are approximately equal or approximately equal up to a multiplication by a small integer power of 2. We did not yet pursue such optimization since we believe the improvements will be small (by a few percent)]. Such transformation can be accomplished directly at the cost of $m \cdot \text{Cost}(R_z(\theta))$ τ gates, but a better approach is to induce this set of gates via the calculation of the input weight sum and the application of $R_z(2^{\lfloor \log(m) \rfloor} \theta)$, $R_z(2^{\lfloor \log(m) \rfloor - 1} \theta)$, ..., $R_z(2\theta)$, and $R_z(\theta)$, to the binary digits (qubits) of the integer number $(\lfloor \log(m) \rfloor + 1)$ -qubit ket representing the input weight of the m qubits needing the application of $R_z(\theta)$ gates, at the cost of $4(m - \text{Weight}(m)) + (\lfloor \log(m) \rfloor + 1) \cdot \text{Cost}(R_z(\theta))$ (where $\text{Weight}(m)$ computes the number of ones in the binary expansion of the integer number m) τ gates. Indeed, for the implementations we considered $\text{Cost}(R_z(\theta)) \approx 50$ τ gates, and thus the saving in the τ count is substantial. To induce the input weight calculation at the cost of at most $4(m - \text{Weight}(m))$ gates, we use $m - \text{Weight}(m)$ full and half adders. We perform all lower digit summations first, using full adders as much as possible (it is always possible when there are at least three bits to be added) and while doing so in parallel. Since both half and full adders need one relative phase Toffoli gate to be computed, costing 4 τ gates each, and only Clifford gates and one measurement to be uncomputed,³¹ the τ count in the calculation of the input weight and proper reset of ancillae is precisely $4(m - \text{Weight}(m))$. The overall optimization of the τ count from applying parallel R_z gates through the calculation of the input weight ranges from 51 to 60%, being better for simulations of higher degree graphs. Indeed, for higher degree graphs, the product formula algorithm expels a higher fraction of resources on implementing Hamiltonian interactions compared with the resources spent on implementing the random disorder.

We use gridsynth^{32,33} to synthesize optimal single-qubit Clifford + τ circuits implementing individual R_z gates. A better strategy relies on the RUS circuits³⁴ that use ancilla, measurement, and feedforward to reduce

Fig. 6 Parameter r scaling data for pre-fault tolerant implementation as a function of the system size n for 4th (left) and 6th (right) order formulas. The black squares, orange circles, blue triangles, and red circles denote $k = 3, 4, 5$, and 7 graphs, respectively. The error bars denote one standard deviation. The solid lines are the best fit power law curves (4) and (5) for 4th and 6th orders, respectively

the τ count in the implementation of a single R_z gate by a factor of 2.5 on average. However, a fully automated implementation of the RUS strategy is unavailable, and the implementation we have is labor-intensive. Thus, we do not compute RUS circuits explicitly, but rather report the respective expected gate counts. The RUS implementation can be easily included into our software as an external and independent package. The result of the overall optimization of the Hamiltonian dynamics simulation circuits by applying the RUS approach ranges between 44 and 50%, with optimization quality favoring smaller order graphs. Indeed, those have smaller parts composed with half and full adders that are not optimized by the RUS.

We note that optimizations described in the last two paragraphs reduce the τ count at the cost of introducing a number of $CNOT$ gates. We keep

track of the $CNOT$ gates to make sure their cost does not overwhelm that of the τ gates.

Empirical bound. To determine the empirical bound on the number r of iterations of the product formula for our problem that aims to address the graphs (3–5–70), (4–4–98), (5–3–72), and (7–2–50), we considered random regular graphs with $k=3, 4, 5$, and 7 , generated using the random matching approach,³⁵ where the number of vertices n ranges from $k+1$ to 12 . For the cases when nk is odd and it is impossible to generate the corresponding random k -regular graph (the number of edge ends can only be even), we take a degree- k random graph with $n+1$ vertices, and remove a randomly selected vertex as well as all edges leading to it. We

Fig. 7 Parameter r scaling data for fault-tolerant implementation as a function of the system size n for 4th (left) and 6th (right) order formulas. The black squares, orange circles, blue triangles, and red circles denote $k=3, 4, 5$, and 7 graphs, respectively. The error bars denote one standard deviation. The solid lines are the best fit power law curves (6) and (7) for 4th and 6th orders, respectively

Table 2. Gate counts for different graphs using 4th order PF

Graph ($k-d-n$)	Pre-fault tolerant cost			Fault-tolerant cost (Gridsynth)			Fault-tolerant cost (RUS)		
	n	$CNOT$	R_z	n	$CNOT$	T	n	$CNOT$	T
(3–1–4)	4	21,420	22,854	6	44,546	656,999	6	312,218	288,084
(3–d–5)	5	25,740	28,606	7	56,538	1,092,171	7	404,791	373,904
(3–d–6)	6	42,030	45,774	10	105,944	1,354,329	12	654,011	595,619
(3–d–7)	7	34,527	37,921	11	109,184	1,053,768	13	533,148	470,965
(3–d–8)	8	49,014	52,915	15	190,432	1,402,834	18	743,399	640,570
(3–d–9)	9	61,776	67,404	16	272,375	1,874,968	19	1,003,682	865,932
(3–d–10)	10	65,850	71,132	18	282,579	1,970,933	21	1,053,528	906,787
(3–d–11)	11	68,364	74,287	19	297,681	1,950,765	22	1,058,286	897,899
(3–d–12)	12	81,840	88,896	22	379,351	2,151,939	25	1,213,604	1,001,577
(4–1–5)	5	31,752	33,523	7	70,338	973,842	7	464,312	428,246
(4–d–6)	6	33,894	36,054	10	101,048	917,418	12	466,553	411,243
(4–d–7)	7	47,475	50,649	11	174,643	1,154,595	13	626,252	528,012
(4–d–8)	8	56,952	60,353	15	231,522	1,493,720	18	812,052	690,186
(4–d–9)	9	70,848	75,288	16	313,316	1,832,632	19	1,019,675	857,497
(4–d–10)	10	74,445	78,713	18	330,598	1,928,493	21	1,074,228	902,955
(4–d–11)	11	88,101	93,387	19	399,603	2,128,496	22	1,217,455	1,000,280
(4–d–12)	12	89,964	95,692	22	434,945	2,220,409	25	1,282,712	1,052,963
(5–1–6)	6	37,584	39,448	10	119,778	930,247	12	486,559	421,418
(5–d–7)	7	37,758	39,797	11	143,125	866,847	13	479,031	399,411
(5–d–8)	8	66,636	69,732	15	275,838	1,627,403	18	902,075	760,854
(5–d–9)	9	62,520	65,658	16	269,112	1,551,995	19	861,283	730,606
(5–d–10)	10	87,354	91,166	18	397,923	2,036,766	21	1,174,295	966,111
(5–d–11)	11	89,817	94,109	19	401,968	2,179,369	22	1,234,464	1,029,598
(5–d–12)	12	109,836	115,276	22	541,352	2,445,309	25	1,462,680	1,177,035
(7–1–8)	8	76,440	78,412	15	362,435	1,814,698	18	1,046,652	873,457
(7–d–9)	9	63,510	65,712	16	296,818	1,482,322	19	854,973	711,118
(7–d–10)	10	72,072	74,633	18	340,054	1,494,097	21	899,034	723,474
(7–d–11)	11	83,283	85,645	19	402,510	1,729,234	22	1,048,114	838,359
(7–d–12)	12	105,534	109,158	22	539,611	2,026,914	25	1,288,877	996,683

Table 3. Gate counts for different graphs using 6th order PF

Graph ($k-d-n$)	Pre-fault tolerant cost			Fault-tolerant cost (Gridsynth)			Fault-tolerant cost (RUS)		
	n	CNOT	R_z	n	CNOT	T	n	CNOT	T
(3-1-4)	4	19,050	20,326	6	40,104	559,199	6	267,051	245,811
(3-d-5)	5	27,900	31,006	7	59,300	851,919	7	403,992	372,413
(3-d-6)	6	43,200	47,048	10	105,779	1,265,212	12	615,321	557,897
(3-d-7)	7	44,370	48,729	11	127,795	1,162,152	13	591,622	519,905
(3-d-8)	8	61,740	66,651	15	216,250	1,512,291	18	808,460	693,154
(3-d-9)	9	73,260	79,932	16	292,799	1,921,586	19	1,039,192	892,361
(3-d-10)	10	87,000	93,974	18	335,036	2,227,206	21	1,202,108	1,030,323
(3-d-11)	11	98,415	106,935	19	382,947	2,408,245	22	1,320,426	1,116,307
(3-d-12)	12	114,855	124,751	22	486,069	2,625,050	25	1,501,123	1,231,756
(4-1-5)	5	30,780	32,497	7	64,205	842,883	7	404,922	372,437
(4-d-6)	6	44,730	47,578	10	117,715	1,051,819	12	534,980	471,445
(4-d-7)	7	57,375	61,209	11	191,469	1,221,648	13	667,424	560,700
(4-d-8)	8	69,720	73,881	15	264,421	1,655,404	18	906,203	768,382
(4-d-9)	9	93,120	98,952	16	359,050	2,011,995	19	1,130,662	946,685
(4-d-10)	10	107,100	113,234	18	428,640	2,415,220	21	1,357,074	1,136,837
(4-d-11)	11	130,455	138,275	19	534,838	2,764,991	22	1,595,466	1,306,936
(4-d-12)	12	141,120	150,096	22	623,051	3,096,764	25	1,803,268	1,477,260
(5-1-6)	6	44,550	46,758	10	129,405	982,123	12	516,259	446,710
(5-d-7)	7	54,405	57,339	11	187,022	1,111,146	13	618,576	515,215
(5-d-8)	8	72,360	75,721	15	294,048	1,684,602	18	939,612	790,538
(5-d-9)	9	82,800	86,952	16	328,222	1,841,418	19	1,029,388	870,758
(5-d-10)	10	108,330	113,054	19	446,178	2,218,756	21	1,289,821	1,057,906
(5-d-11)	11	132,000	138,615	20	521,262	2,794,246	22	1,591,495	1,327,663
(5-d-12)	12	157,140	164,916	22	699,399	3,077,140	25	1,858,714	1,491,134
(7-1-8)	8	76,440	78,412	15	332,706	1,607,197	18	935,256	776,724
(7-d-9)	9	89,610	92,712	16	365,610	1,781,648	19	1,035,870	859,263
(7-d-10)	10	99,990	103,539	19	436,370	1,879,136	21	1,135,877	911,157
(7-d-11)	11	115,020	118,276	20	507,581	2,177,549	22	1,321,786	1,058,741
(7-d-12)	12	141,570	146,426	22	646,853	2,360,222	25	1,513,669	1,165,082

then insert edges connecting $\lfloor k/2 \rfloor$ non-overlapping pairs of the resulting k degree- $(k-1)$ vertices, chosen at random, provided that the introduction of the new edges does not lead to a multi-edged graph.

Figures 6 and 7 show, for pre-fault tolerant and fault-tolerant implementations, respectively, the r scaling for the 4th and 6th order formulas for degree $k=3, 4, 5$, and 7 random regular graphs. By performing least square linear fitting on the log-log scale, we determined the following scaling of r in the pre-fault tolerant case, 4th order formula

$$r_3 = 116.3n^{0.169}, r_4 = 66.4n^{0.331}, r_5 = 30.1n^{0.602}, \quad (4)$$

in the pre-fault tolerant case, 6th order formula

$$r_3 = 12.5n^{0.564}, r_4 = 7.05n^{0.759}, r_5 = 4.24n^{0.883}, r_7 = 7.11n^{0.476}, \quad (5)$$

while for the fault-tolerant case, 4th order formula

$$r_3 = 126n^{0.215}, r_4 = 80.5n^{0.328}, r_5 = 34.6n^{0.620}, \quad (6)$$

and for the fault-tolerant case, 6th order formula

$$r_3 = 15.9n^{0.493}, r_4 = 7.82n^{0.751}, r_5 = 5.26n^{0.826}, r_7 = 7.61n^{0.490}, \quad (7)$$

where r_k is the empirical bound for the degree- k graph with evolution time $t = 2d = 10, 8, 6$, and 4 for graphs with $k = 3, 4, 5$, and 7, respectively. We did not include the scaling for the 4th order formula and $k = 7$, since there were too few points to study and there did not seem to be enough stability in the data. However, we investigated different values ϵ and empirically confirmed the stability in the scalings for $k = 3, 4$, and 5.

The detailed results showing concrete gate counts for individual cases are available in Tables 2 and 3 for the 4th and 6th order formulas, respectively. Also shown are the gate counts expected for RUS approach,³⁴

where, for each r_z approximation, we expect a factor 2.5 reduction in the τ count at the cost of the addition of at most $(\tau\text{-count} + 1)$ CNOT gates and one ancilla.

DATA AVAILABILITY

The benchmark circuits considered in this paper are available at https://github.com/y-nam/Heisenberg_Interaction.

CODE AVAILABILITY

The software implementation used to produce the optimized circuits is not publicly available.

ACKNOWLEDGEMENTS

Authors thank Prof. Scott Aaronson (University of Texas, Austin) and Prof. N. Julien Ross (Dalhousie University) for helpful discussions. This material was based on work supported by the National Science Foundation, while DM working at the Foundation. Any opinion, finding, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

AUTHOR CONTRIBUTIONS

All authors researched, collated, and wrote this paper.

ADDITIONAL INFORMATION

Competing interests: Provisional patent applications for this work were filed by IonQ, Inc.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

REFERENCES

- Boxio, S. et al. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600 (2018).
- Reiher, M., Wiebe, N., Svore, K. M., Wecker, D. & Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proc. Natl Acad. Sci.* **114**, 7555–7560 (2017).
- Kutin, S. A. Shor's algorithm on a nearest-neighbor machine. Preprint at <https://arxiv.org/abs/quant-ph/0609001> (2006).
- Babbush, R. et al. Encoding electronic spectra in quantum circuits with linear T complexity. *Phys. Rev. X* **8**, 041015 (2018).
- Childs, A. M., Maslov, D., Nam, Y. S., Ross, N. J. & Su, Y. Toward the first quantum simulation with quantum speedup. *Proc. Natl Acad. Sci.* **115**, 9456–9461 (2018).
- Wang, Y. et al. Single-qubit quantum memory exceeding ten-minute coherence time. *Nat. Photonics* **11**, 646–650 (2017).
- Linke, N. M. et al. Experimental comparison of two quantum computing architectures. *Proc. Natl Acad. Sci.* **114**, 3305–3310 (2017).
- Campbell, E. Shorter gate sequences for quantum computing by mixing unitaries. *Phys. Rev. A* **95**, 042306 (2017).
- Gaebler, J. P. et al. High-fidelity universal gate set for $^9\text{Be}^+$ ion qubits. *Phys. Rev. Lett.* **117**, 060505 (2016).
- Zhang, J. et al. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature* **551**, 601–604 (2017).
- Bernien, H. et al. Probing many-body dynamics on a 51-atom quantum simulator. *Nature* **551**, 579–584 (2017).
- Leung, P. H. & Brown, K. R. Entangling an arbitrary pair of qubits in a long ion crystal. *Phys. Rev. A* **98**, 032318 (2018).
- Luitz, D. J., Laflorencie, N. & Alet, F. Many-body localization edge in the random-field Heisenberg chain. *Phys. Rev. B* **91**, 081103 (2015).
- Nandkishore, R. & Huse, D. A. Many body localization and thermalization in quantum statistical mechanics. *Annu. Rev. Condens. Matter Phys.* **6**, 15–38 (2015).
- Huse, D. A. & Pal, A. The many-body localization phase transition. *Phys. Rev. B* **82**, 174411 (2010).
- Chen, J. et al. Classical simulation of intermediate-size quantum circuits. Preprint at <https://arxiv.org/abs/1805.01450> (2018).
- Li, R., Wu, B., Ying, M., Sun, X., & Yang, G. Quantum supremacy circuit simulation on Sunway TaihuLight. Preprint at <https://arxiv.org/abs/1804.04797> (2018).
- Smelyanskiy, M., Sawaya, N. P. D. & Aspuru-Guzik, A. qHiPSTER: The quantum high performance software testing environment. Preprint at <https://arxiv.org/abs/1601.07195> (2018).
- Lieb, E. & Robinson, D. The finite group velocity of quantum spin systems. *Commun. Math. Phys.* **28**, 251–257 (1972).
- Suzuki, M. General theory of fractal path integrals with applications to many-body theories and statistical physics. *J. Math. Phys.* **32**, 400–407 (1991).
- Berry, D. W., Ahokas, G., Cleve, R. & Sanders, B. C. Efficient quantum algorithms for simulating sparse Hamiltonians. *Commun. Math. Phys.* **270**, 359–371 (2007).
- Leung, P. H. et al. Robust two-qubit gates in a linear ion crystal using a frequency-modulated driving force. *Phys. Rev. Lett.* **120**, 020501 (2018).
- Alegre, I., Fiol, M. A. & Yebra, J. L. A. Some large graphs with given degree and diameter. *J. Graph Theory* **10**, 219–224 (1986).
- Exoo, G. Large regular graphs of given degree and diameter. <http://isu.indstate.edu/ge/DD/index.html>. (last accessed 19 January 2018).
- Hoffman, A. J. & Singleton, R. R. Moore graphs with diameter 2 and 3. *IBM J. Res. Dev.* **5**, 497–504 (1960).
- Wikipedia. Table of the largest known graphs of a given diameter and maximal degree. https://en.wikipedia.org/wiki/Table_of_the_largest_known_graphs_of_a_given_diameter_and_maximal_degree. (last accessed 4 May 2018).
- Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. In *Proc. 43rd Annual ACM Symposium on Theory of Computing (STOC'11)* 333–342, (San Jose, CA, USA, 2011).
- Januario, T. & Urrutia, S. An edge coloring heuristic based on Vizing's theorem. In *Proc. Proceedings of the Brazilian Symposium on Operations Research*, 3994–4002, (Rio de Janeiro, Brazil, 2012).
- Holyer, I. The NP-completeness of edge-coloring. *SIAM J. Comput.* **10**, 718–720 (1981).
- Nam, Y., Ross, N. J., Su, Y., Childs, A. M. & Maslov, D. Automated optimization of large quantum circuits with continuous parameters. *npj: Quantum Inf.* **4**, 23 (2018).
- Gidney, C. Halving the cost of quantum addition. *Quantum* **2**, 74 (2018).
- Kliuchnikov, V., Maslov, D. & Mosca, M. Fast and efficient exact synthesis of single-qubit unitaries generated by Clifford and τ gates. *Quantum Inf. Comput.* **13**, 607–630 (2013).
- Ross, N. J. & Selinger, P. Optimal ancilla-free Clifford + T approximation of z-rotations. *Quantum Inf. Comput.* **16**, 901–953 (2016).
- Bocharov, A., Roetteler, M. & Svore, K. M. Efficient synthesis of universal Repeat-Until-Success circuits. *Phys. Rev. Lett.* **114**, 080502 (2015).
- Wormald, N. Models of Random Regular Graphs. *Surveys in Combinatorics*, 239–298, (Cambridge University Press, Cambridge, UK, 1999).
- Vatan, F. & Williams, C. Optimal quantum circuits for general two-qubit gates. *Phys. Rev. A* **69**, 032315 (2004).
- Nam, Y., Su, Y. & Maslov, D. Approximate quantum fourier transform with $O(n \log(n))$ τ gates. Preprint at <https://arxiv.org/abs/1803.04933> (2018).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019