

Implementing a fast unbounded quantum fanout gate using power-law interactions

Andrew Y. Guo^{1,2,*} Abhinav Deshpande^{1,2} Su-Kuan Chu^{1,2} Zachary Eldredge^{1,2} Przemyslaw Bienias^{1,2}
Dhruv Devulapalli^{1,2} Yuan Su^{1,3,4} Andrew M. Childs^{1,3,4} and Alexey V. Gorshkov^{1,2}

¹Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, Maryland 20742, USA

²Joint Quantum Institute, NIST/University of Maryland, College Park, Maryland 20742, USA

³Department of Computer Science, University of Maryland, College Park, Maryland 20742, USA

⁴Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland 20742, USA



(Received 12 June 2022; accepted 12 October 2022; published 27 October 2022)

The standard circuit model for quantum computation presumes the ability to directly perform gates between arbitrary pairs of qubits, which is unlikely to be practical for large-scale experiments. Power-law interactions with strength decaying as $1/r^\alpha$ in the distance r provide an experimentally realizable resource for information processing, whilst still retaining long-range connectivity. We leverage the power of these interactions to implement a fast quantum fanout gate with an arbitrary number of targets. Our implementation allows the quantum Fourier transform (QFT) and Shor's algorithm to be performed on a D -dimensional lattice of qubits in time that scales (poly)logarithmically in the number of qubits, using interactions with $\alpha \leq 2D$. As a corollary, we show that power-law systems with $\alpha \leq 2D$ are difficult to simulate classically even for short times, under a standard assumption that factoring is classically intractable. Complementarily, we develop a technique to give a general lower bound—linear in the size of the system—on the time required to implement the QFT and the fanout gate in systems that are constrained by a linear light cone. This allows us to prove an asymptotically tighter lower bound for long-range systems than was possible with previously available techniques.

DOI: [10.1103/PhysRevResearch.4.L042016](https://doi.org/10.1103/PhysRevResearch.4.L042016)

In the circuit model for quantum computation, the depth of a quantum circuit is given by the number of layers of nonoverlapping quantum gates it contains. In typical quantum systems, coherence times are a limitation, so low-depth circuits prioritized for the regime of noisy intermediate-scale quantum computers are more desirable [1]. Various proposed models of quantum computation are equivalent up to polynomial overhead, making the definition of the complexity class BQP insensitive to the model of computation [2–5].

However, these models can differ in the precise complexity of operations. As a drastic example, suppose we are given access to a fast unbounded fanout gate represented by the map $|x\rangle|y_1\rangle|y_2\rangle \dots \mapsto |x\rangle|y_1 \oplus x\rangle|y_2 \oplus x\rangle \dots$ where the $\oplus$ operator denotes bitwise XOR (bit y_i is flipped if $x = 1$ and not flipped otherwise). This operation is a reversible analog of a gate that copies x to registers $y_1, y_2, \dots$. By “unbounded,” we mean that there is no limit on the number of bits that can be targeted by this operation.

The unbounded fanout gate makes it possible for constant-depth circuits to perform a number of fundamental quantum arithmetic operations [5]. Furthermore, unbounded fanout can also reduce the quantum Fourier transform (QFT)—a subrou-

tine of a large class of quantum algorithms, including most famously Shor's algorithm for integer factorization [6]—to constant depth as well. In fact, it enables implementing the entirety of Shor's algorithm by constant-depth circuits with access to a polynomial amount of classical pre- and postprocessing [7].

While the unbounded fanout gate is clearly a powerful resource for quantum computation, its efficient implementation in physically realizable architectures has not been studied in great depth. In the standard circuit model—where one may apply single-qubit and two-qubit gates from a standard gate set on arbitrary nonoverlapping subsets of the qubits—a fanout gate on n qubits can be implemented optimally in $\Theta(\log n)$ -depth [8,9]. One may also consider the Hamiltonian model, in which one may apply single-qubit and two-qubit Hamiltonian terms. In particular, in the Hamiltonian model with all-to-all unit-strength interactions, one can implement the fanout gate in constant time [10,11]. However, the assumption of being able to directly apply interactions between two arbitrarily distant qubits does not hold in practice for large quantum computing architectures [12–15]. Mapping these circuits to restricted architectures inevitably leads to overheads and potentially even different asymptotic scaling. In D -dimensional nearest-neighbor architectures, for example, the unbounded fanout gate can only be implemented unitarily in depth $\Theta(n^{1/D})$ [16]. And while there exist protocols that can implement the fanout gate in constant depth on these architectures [17], these proposals require intermediate measurements along with classical control—a resource that may be inaccessible in certain near-term experimental systems

*guoa@umd.edu

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

[18]. The overheads resulting from such physical restrictions could therefore limit the potential asymptotic speed-up from a fast quantum fanout.

Systems with power-law interactions, however, present an opportunity for realizing these speed-ups. Specifically, for a lattice of qubits in D dimensions, the interaction strengths between pairs of qubits separated by a distance r are weighted by a power-law decaying function $1/r^\alpha$. These long-range interactions are native to many experimental quantum systems and have attracted interest as potential resources for faster quantum information processing. Examples of long-range interactions include dipole-dipole and van der Waals interactions between Rydberg atoms [19,20], and dipole-dipole interactions between polar molecules [21] and between defect centers in diamond [20,22]. Previous papers have explored the acceleration of quantum information processing using strong and tunable power-law interactions between Rydberg states [23–29], which can implement k -local gates that control or target simultaneously $k \gg 10$ qubits. Those gates still have a finite spatial range and can therefore only give a constant-factor speed-up over nearest-neighbor architectures. Recently, Refs. [30–33] gave protocols that take advantage of power-law interactions to quickly transfer a quantum state across a lattice. As we will show, it is also possible to leverage the power of these interactions to implement quantum gates asymptotically faster than is possible with finite-range interactions [34].

In this Letter, we propose a protocol for implementing the unbounded fanout gate quickly using engineered Hamiltonians with power-law interactions. As an application of this protocol, we show that simulating strongly long-range systems with $\alpha \leq 2D$ for logarithmic time or longer is classically intractable, if factoring is classically hard. As a complement to our upper bounds on the fanout time, we also develop a technique that allows us to prove the tightest-known lower bounds for the time required to implement the QFT and unbounded fanout in a general lattice architecture.

Protocol for fast fanout using long-range interactions. To perform a fanout gate on n logical qubits, we employ as subroutines modified versions of the state transfer protocols from Refs. [30] and [35] to generate a many-body entangled state in $\mathcal{O}(\text{polylog}(n))$ time using long-range interactions with $\alpha < 2D$.

As an intermediate step, both state-transfer protocols “broadcast” a single-qubit state into the corresponding GreenbergerHorneZeilinger (GHZ)-like state,

$$(\psi_0|0\rangle + \psi_1|1\rangle) \otimes |00\dots 0\rangle \mapsto \psi_0|00\dots 0\rangle + \psi_1|11\dots 1\rangle, \quad (1)$$

where $\psi_0, \psi_1 \in \mathbb{C}$ and $|\psi_0|^2 + |\psi_1|^2 = 1$. In the case of Ref. [30], this long-range broadcast is achieved by performing a sequence of cascaded controlled-NOT (CNOT) gates—similar to the standard gate-based implementation of the unbounded fanout gate. The CNOT gate from qubit i to qubit j can be implemented by a Hamiltonian $H_{ij} = h_{ij}|1\rangle\langle 1|_i \otimes X_j$ acting for time $t = \pi/(2h_{ij})$, up to a local unitary. Applying a Hamiltonian $H(t) = \sum_{ij} H_{ij}(t)$, which variously turns on/off interactions between pairs of qubits at different times, allows one to implement the broadcast in Eq. (1).

By using Hamiltonians with long-range interactions h_{ij} satisfying $||h_{ij}|| \leq 1/r_{ij}^\alpha$, it is possible to implement the

Algorithm 1. Implementing fanout with long-range interactions.

```

1: Initialize ancillary qubits:  $|a_i\rangle \leftarrow |0\rangle$  for  $i = 1$  to  $n$ 
2: CNOT( $|d_1\rangle \rightarrow |a_1\rangle$ )
    $\nabla$  Apply broadcast operation as shown in Eq. (1) to ancillae:
3: LONGRANGEBROADCAST( $|a_1\rangle \rightarrow |a_2\rangle, \dots, |a_n\rangle$ )
    $\nabla$  parfor indicates that for-loop can be implemented in parallel
4: parfor  $i = 2$  to  $n$  do
5:   CNOT( $|a_i\rangle \rightarrow |d_i\rangle$ )  $\triangleright$  Transfer fanout to data qubits
6: end parfor
    $\nabla$  Apply broadcast operation in reverse to uncompute ancillae
7: REVERSELONGRANGEBROADCAST( $|a_2\rangle, \dots, |a_n\rangle \rightarrow |a_1\rangle$ )
8: CNOT( $|d_1\rangle \rightarrow |a_1\rangle$ )

```

broadcast operation asymptotically faster than with short-range interactions—a statement that we will prove rigorously later in the text. For a system of n qubits, Refs. [30] and [35] showed that this broadcast time depends on the power-law exponent α and the dimension of the system D as follows:

$$t_{\text{GHZ}} = \begin{cases} \mathcal{O}(n^0) & \alpha < D \\ \mathcal{O}(\log n) & \alpha = D \\ \mathcal{O}(\log^{\kappa_\alpha} n) & \alpha \in (D, 2D) \\ \mathcal{O}(e^{\gamma\sqrt{\log n}}) & \alpha = 2D \\ \mathcal{O}(n^{\min(\alpha/D-2, 1)}) & \alpha > 2D, \end{cases} \quad (2)$$

for constants $\gamma = 3\sqrt{d}$ and $\kappa_\alpha = \log 4 / \log(2D/\alpha)$. We term the broadcast time t_{GHZ} , since it corresponds to the GHZ-state-construction time when $\psi_0 = \psi_1 = 1/\sqrt{2}$. This long-range broadcast is not the same as fanout because it requires that all intermediary qubits (besides the first qubit) be initialized in the $|0\rangle$ state. However, as we now show, it is possible to adapt the broadcast protocol to implement the fanout gate in time t_{GHZ} using n ancillary qubits.

Consider a system of n data qubits arranged on a D -dimensional lattice. Furthermore, assume there are n ancillary qubits, each located adjacent to one of the original qubits. We denote the qubits as $|d_1\rangle, |d_2\rangle, \dots, |d_n\rangle$ and $|a_1\rangle, |a_2\rangle, \dots, |a_n\rangle$ for data and ancilla, respectively. Suppose we want to perform fanout with $|d_1\rangle$ as control, and that all ancillae are guaranteed to be in state $|0\rangle$. Then the sequence of operations in Algorithm 1 (also depicted graphically in Fig. 1) implements the fanout operation.

In addition to accomplishing fanout, this protocol returns the ancillary qubits to the $|0\rangle$ state. Modulo the $\mathcal{O}(n)$ short-range operations that can be done in parallel in a single time step, the protocol requires time $2t_{\text{GHZ}}$. Hence, it can implement the fanout gate in time that is constant for $\alpha < D$, polylogarithmic for $D \leq \alpha < 2D$, and polynomial for $\alpha > 2D$.

We briefly comment on the constant-depth implementation of the QFT and Shor’s algorithm using the unbounded fanout gate. An n -qubit QFT circuit can be performed with $\mathcal{O}(n \log n)$ gates to $1/\text{poly}(n)$ precision [36]. Using unbounded fanout, the circuit can be reduced to constant depth with $\mathcal{O}(n \log n)$ ancillary qubits [5]. We note that including these ancillae in the lattice would not change the asymptotic scaling of our protocol for $\alpha < 2D$, since t_{GHZ} is $\mathcal{O}(\text{polylog}(n))$ in this regime.




FIG. 1. A protocol for a fast unbounded quantum fanout gate using long-range interactions, depicted here for a 1D lattice. The layout consists of a chain of data qubits, along with their adjacent ancillary qubits that are initialized to $|0\rangle$. (a) The first step is a local controlled-NOT (CNOT) gate from $|d_1\rangle$ to $|a_1\rangle$. (b) The application of the long-range “broadcast” from $|a_1\rangle$ to the rest of the ancillary qubits $|a_i\rangle$ creates a GHZ-like state in Eq. (1) for the ancillary qubits together with the first data qubit. (c) We apply CNOT gates from ancillary qubit $|a_i\rangle$ to the data qubit $|d_i\rangle$, which can be done in parallel. After this step, we reverse process (b) and process (a) to return the ancillary qubits to $|0\rangle$ (not redrawn here).

We also remark on the potential ways of performing the fast fanout gate experimentally. One method would be to use two different qubit realizations for the data and ancillary qubits, with the former interfacing with the latter through local interactions. For example, one may consider a system of fermionic alkaline-earth atoms, with each atom’s electron on the clock transition acting as an ancillary qubit, its nuclear spin as the data qubit, and with long-range Rydberg-Rydberg interactions between electrons. In this heterogeneous case, the n CNOT gates between the data and ancilla qubits of our protocol would be straightforward to implement in parallel [37]. Alternatively, one may consider implementing the protocol entirely in a system of Rydberg atoms using power-law interactions of different strengths. This could be done using a gate based on van der Waals interactions ($\alpha = 6$) to implement the short-range CNOT gates between data and ancilla qubits, and then rotating into a new basis in order to implement the long-range broadcast interaction using dipole-dipole interactions

($\alpha = 3$). We provide more discussion on the implementation of the protocol in the Supplemental Material [38].

Intractability of classical simulation of strongly long-range systems. As a corollary, the protocol shows that strongly long-range interacting systems with $\alpha < 2D$ evolving for time polylogarithmic in n or longer are difficult to simulate classically in the worst case. The argument operates by a complexity-theoretic reduction from integer factoring, a problem that is assumed to be difficult for classical computers with the ability to use random bits (FACTORING $\notin$ BPP). This assumption suggests that the simulation of strongly long-range systems could provide an avenue towards a potentially useful experimental demonstration of quantum computational advantage.

The argument proceeds as follows. The time required to implement the fanout gate using Algorithm 1 is $\mathcal{O}(\text{polylog}(n))$ for $\alpha < 2D$. It is possible to implement Shor’s order-finding algorithm in time $\mathcal{O}(t_{\text{FO}})$ using a small amount of classical preprocessing (polynomial in n) [5,39]. Using the ability to sample from the output of the order-finding algorithm to error $\varepsilon < 0.4 < 4/\pi^2$, classically efficient postprocessing can output a factor of an n -bit integer with probability $\Omega(1)$ [6]. Therefore, if it were possible to efficiently sample from the output distribution in strongly long-range systems for evolution-time $t = \mathcal{O}(\log n)$, then it would be possible to factor n -bit integers efficiently as well. The best classical algorithm currently known for factoring an n -bit integer takes runtime $\exp[\mathcal{O}(\sqrt{n \log n})]$ [40] and the problem is widely believed to be classically intractable. This stands in contrast to systems with finite-range interactions in 1D, for which efficient classical simulation is possible up to any time satisfying $t \leq \mathcal{O}(\log n)$ [41]. Under the complexity assumption mentioned above (FACTORING $\notin$ BPP), we have shown that this result is not fully generalizable to strongly long-range interacting systems [42].

Lower bounds on the time required to implement QFT and fanout. In previous sections, we found upper bounds to the time required to implement fanout—and by corollary, the QFT. As a way to benchmark our long-range protocol, we now proceed to discuss lower bounds for implementing fanout and the QFT. Recall that the protocol in Algorithm 1 can implement fanout in time t_{GHZ} , which scales as $\mathcal{O}(\text{polylog}(n))$ for long-range systems with $\alpha < 2D$. In this section, we show that such fast asymptotic runtimes cannot be achieved in architectures with strict locality constraints.

In Ref. [43], Maslov showed that a specific way of implementing the QFT requires $\Omega(n)$ depth on the 1D nearest-neighbor architecture, although this does not rule out other QFT implementations with potentially sublinear depth. Here, we devise a technique that yields a lower bound of $\Omega(n^{1/D})$ for the time required to perform a QFT in the Hamiltonian model. This result strengthens and generalizes Maslov’s bound to higher dimensions and to the Hamiltonian model. In addition, we show that the same lower bound applies even to circuits that perform the QFT approximately.

Our $\Omega(n^{1/D})$ lower bound holds for any lattice system with finite velocities of information spreading, which include short-range interactions (i.e., finite-ranged or exponentially decaying) and power-law interactions with $\alpha > 2$ in one dimension or $\alpha > 2D + 1$ for $D > 1$ [33,44,45]. Combined

with our results above, this implies that systems with strongly long-range interactions can implement the QFT and fanout asymptotically faster than more weakly interacting systems.

The intuitive idea behind our proof is that the QFT unitary can spread out operators in a certain precise sense, a task that can be bounded by the “Frobenius-norm light cone” of Ref. [32]. The fact that this light cone imposes a finite speed limit on information propagation in short-range interacting systems implies that the minimum time $t_2(r)$ required for operator-spreading is proportional to the distance between qubits r . This constrains the implementation time for the QFT, denoted t_{QFT} , by $\Omega(n^{1/D})$, from which the circuit-depth lower bound follows. The same Frobenius-norm bound also constrains the time required to implement the approximate QFT (AQFT).

We consider the 4^n -dimensional vector space of n -qubit operators for which the set of Pauli operators $\{I, X, Y, Z\}^{\otimes n}$ forms a basis. We quantify operator spreading outside a region of radius r as follows. Taking an operator $|O\rangle$ initially supported on site 1, we measure the weight of its time-evolved version, $|O(t)\rangle$, on sites at distance r (and beyond) using a projection operator $\mathcal{Q}_r$, which projects onto strings of Pauli operators that act nontrivially on at least one site at distance r or greater. We measure the weight of this projected operator $|O_r\rangle := \mathcal{Q}_r|O(t)\rangle$ via the (squared) normalized Frobenius norm $\|O_r\|_F^2 := \text{Tr}(O_r^\dagger O_r)/2^n$, which coincides with the Euclidean norm over the operator space, $(O_r|O_r)$ [46]. We define $t_2(r)$ to be the minimal time after which $\|O_r\|_F^2$ is able to achieve a predetermined constant value [32].

We show that operators spread by the action of the QFT can have high weight on distant regions, which implies that $t_{\text{QFT}} \geq t_2(r)$. Recall that the QFT operator on n qubits is defined as $U_{\text{QFT}} := \sum_{y,z=0}^{2^n-1} |y\rangle\langle z| \omega^{yz} / \sqrt{2^n}$, where $|y\rangle$ ($|z\rangle$) is an n -qubit state that encodes the binary representation of y (z) and $\omega = e^{2\pi i/2^n}$. Then the following lemma holds:

Lemma. Let U_{QFT} be the QFT operator on n qubits arranged in a D -dimensional lattice such that the first and n th qubits are located a distance $r = \Theta(n^{1/D})$ apart. Then $U_{\text{QFT}}^\dagger Z_1 U_{\text{QFT}} =: Z_1'$ is an operator with at least constant weight at distance r .

We defer the (short) proof of this Lemma to the Supplemental Material [38]. As a result of the Lemma, t_{QFT} follows the light cone defined by the normalized Frobenius norm, which is at least as stringent as the Lieb-Robinson light cone. This leads directly to the following theorem:

Theorem. For systems with finite-range or exponentially-decaying interactions in D dimensions, the time required to implement the QFT unitary is lower bounded by $t_{\text{QFT}} = \Omega(r)$, where $r = \Theta(n^{1/D})$ is the distance between the first and n th qubits.

For systems with long-range interactions, the Lieb-Robinson light cone gives the following bounds [33,47–49]:

$$t_{\text{QFT}} = \begin{cases} \Omega(1), & \alpha = D \\ \Omega(\log r), & \alpha \in (D, 2D] \\ \Omega(r^{\alpha-2D}), & \alpha \in (2D, 2D+1] \\ \Omega(r), & \alpha > 2D+1. \end{cases} \quad (3)$$

The Frobenius light cone gives the following (tighter) bounds [50,51]:

$$t_{\text{QFT}} = \begin{cases} \Omega(r^{\frac{2\alpha-2D}{2\alpha-D+1}}), & \alpha \in (D, 2D] \\ \Omega(r^{\alpha-1}), & \alpha \in (1, 2], D = 1 \\ \Omega(r), & \alpha > 2, D = 1. \end{cases} \quad (4)$$

We note that the lower bounds in the Theorem also apply to the fanout time, t_{FO} , through the observation that fanout also performs operator spreading (using X_1 instead of Z_1). We emphasize that these bounds pertain to the Hamiltonian model, where commuting terms can be implemented simultaneously and state transfer could in theory be done in $\mathcal{O}(1)$ time for sufficiently small α .

We observe that the QFT can implement quantum state transfer as well. The goal of state transfer is to find a unitary V such that $V(|\psi\rangle \otimes |0\rangle^{\otimes n-1}) = |0\rangle^{\otimes n-1} \otimes |\psi\rangle$ [30,52]. The unitary $V = H^{\otimes n} U_{\text{QFT}}$ (where H represents the single-qubit Hadamard gate) satisfies this definition of state transfer.

For the AQFT, the lower bound follows in a similar fashion. The circuit that implements the QFT approximately with error ε can be represented by a unitary $\tilde{U}_{\text{QFT}}$ such that $\|U_{\text{QFT}} - \tilde{U}_{\text{QFT}}\| \leq \varepsilon$ [39]. In the Supplemental Material [38], we show that the operator $\tilde{U}_{\text{QFT}}^\dagger Z_1 \tilde{U}_{\text{QFT}}$ also has large support on sites beyond distance r as well, implying that the lower bounds in Eqs. (3) and (4) also hold for the AQFT.

Conclusions and Outlook. In summary, we have developed a fast protocol for the unbounded fanout gate using power-law interactions. For $\alpha \leq 2D$, the protocol can perform the gate parametrically faster than is possible with short-range interactions. In particular, for experimentally realizable dipole-dipole interactions with $\alpha = 3$ in two and three dimensions, as well as van der Waals interactions with $\alpha = 6$ in three dimensions, the fast fanout protocol allows the quantum Fourier transform and Shor’s algorithm to be performed in time that scales subpolynomially in the number of qubits. As a corollary, we showed that the classical simulation of strongly long-range systems with $\alpha < 2D$ for time $t = \mathcal{O}(\text{polylog}(n))$ is at least as difficult as integer factorization, which is believed to be intractable in polynomial time classically.

To benchmark our protocol, we developed a new and general approach for lower-bounding the time required to perform a given quantum algorithm that is independent of its implementation as a quantum circuit. In particular, we derive a $\Omega(n^{1/D})$ lower bound on the time required to implement quantum fanout—as well as the exact and approximate QFTs—for all systems constrained by a linear light cone. In doing so, we used the state-of-the-art Frobenius bounds from [32,50,51], which have been shown to be tighter than the Lieb-Robinson bound in the regimes of α in which they hold. While the Lieb-Robinson bound has been shown to be saturable for all $\alpha \geq D$ (up to logarithmic factors), the corresponding result for the Frobenius bound remains unknown. For higher dimensions, the conjectured Frobenius light cone is $t = \Omega(r^{\alpha-D})$ for $D < \alpha < D+1$ [51]. If this generalization of the Frobenius bound were to hold, our lower bounds on the circuit complexity of the QFT and fanout would immediately generalize.

As a final remark, we have derived our lower bounds on t_{QFT} under the assumption that the first and last qubits of the QFT are separated by a distance of $r = \Theta(n^{1/D})$. However, other mappings of computational qubits to lattice qubits could potentially lead to faster implementations. For example, consider the mapping onto a one-dimensional chain of qubits wherein the second half of the chain is interleaved in reverse order with the first half [53]. Applying the QFT to a product state in this layout results in a state with two-qubit correlations that decay exponentially in the distance between the qubits. In this case, our lower bound techniques cannot rule out the possibility of $t_{\text{QFT}} = o(n)$ for short-range interacting Hamiltonians. This suggests that t_{QFT} could depend strongly on qubit placement. Given that the QFT is typically used as a subroutine for more complex algorithms, it may not always be possible to reassign qubits without incurring costs elsewhere in the circuit. Still, it would be interesting to investigate whether careful qubit placement could yield a faster QFT.

Acknowledgments. We would like to thank Minh Tran, Adam Ehrenberg, Chi-Fang Chen, Andrew Lucas, and Fred Chong for helpful discussions. A.Y.G. would like to thank Dmitri Maslov and Yunseong Nam for informative discussions regarding the QFT. A.Y.G. and D.D. are supported by the NSF Graduate Research Fellowship Program under Grant No. DGE-1840340. We acknowledge support by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Testbed Pathfinder (Award No. DE-SC0019040) and Accelerated Research in Quantum Computing (Award No. DE-SC0020312) programs, the U.S. Army Research Office (MURI Award No. W911NF-16-1-0349), and NSF QLCI (Award No. OMA-2120757). A.Y.G., A.D., S.K.C., Z.E., P.B., and A.V.G. additionally acknowledge support by AFOSR, AFOSR MURI, U.S. Department of Energy Award No. DE-SC0019449, NSF PFCQC program, DoE QSA, ARL CDQI, and NSF PFC at JQI. S.K.C. also acknowledges the support from the Studying Abroad Scholarship by Ministry of Education in Taiwan.

-
- [1] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [2] E. Bernstein and U. Vazirani, Quantum complexity theory, in *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing - STOC '93* (ACM Press, San Diego, 1993), pp. 11–20.
 - [3] R. Raussendorf and H. J. Briegel, A One-Way Quantum Computer, *Phys. Rev. Lett.* **86**, 5188 (2001).
 - [4] R. Raussendorf, D. E. Browne, and H. J. Briegel, Measurement-based quantum computation on cluster states, *Phys. Rev. A* **68**, 022312 (2003).
 - [5] P. Hoyer and R. Spalek, Quantum fanout is powerful, *Theory of Comput.* **1**, 81 (2005).
 - [6] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Comput.* **26**, 1484 (1997).
 - [7] While the rate-limiting step for factoring is actually modular exponentiation, it is possible to implement this subroutine in logarithmic quantum depth with a polynomial number of classical gates as a pre-computation step, using a standard universal gate set for quantum computing [39]. Using the unbounded fanout gate, the quantum circuit for modular exponentiation can be reduced to constant depth as well [5].
 - [8] In this Letter, we use the following standard asymptotic notations: the “big- $\mathcal{O}$ ” notation $f(x) = \mathcal{O}(g(x))$ indicates the existence of constants c and N such that $0 \leq f(x) \leq cg(x)$ for all $x \geq N$; the “big- Ω ” notation $f(x) = \Omega(g(x))$ indicates that $g(x) = \mathcal{O}(f(x))$; the “big- Θ ” notation $f(x) = \Theta(g(x))$ when both $f(x) = \mathcal{O}(g(x))$ and $f(x) = \Omega(g(x))$; finally, the “little- o ” notation $f(x) = o(g(x))$ indicates that $\lim_{x \rightarrow \infty} f(x)/g(x) = 0$. Additionally, we use $\text{poly}(n)$ to refer to computational resources that are polynomial in n .
 - [9] A. Broadbent and E. Kashefi, Parallelizing quantum circuits, *Theor. Comput. Sci.* **410**, 2489 (2009).
 - [10] S. A. Fenner, Implementing the fanout gate by a Hamiltonian, [arXiv:quant-ph/0309163](https://arxiv.org/abs/quant-ph/0309163).
 - [11] S. A. Fenner and Y. Zhang, Implementing fanout, parity, and Mod gates via spin exchange interactions, [arXiv:quant-ph/0407125](https://arxiv.org/abs/quant-ph/0407125).
 - [12] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects, *Phys. Rev. A* **89**, 022317 (2014).
 - [13] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Experimental comparison of two quantum computing architectures, *Proc. Natl. Acad. Sci. USA* **114**, 3305 (2017).
 - [14] A. Bapat, Z. Eldredge, J. R. Garrison, A. Deshpande, F. T. Chong, and A. V. Gorshkov, Unitary entanglement construction in hierarchical networks, *Phys. Rev. A* **98**, 062328 (2018).
 - [15] A. M. Childs, E. Schoute, and C. M. Unsal, Circuit Transformations for Quantum Architectures, in *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)* (Dagstuhl, Germany, 2019), Vol. 135, pp. 3:1–3:24.
 - [16] D. J. Rosenbaum, Optimal quantum circuits for nearest-neighbor architectures, in *8th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2013)* (Dagstuhl, Germany, 2013), Vol. 22, pp. 294–307.
 - [17] P. Pham and K. M. Svore, A 2D nearest-neighbor quantum architecture for factoring in polylogarithmic Depth, *Quantum Inf. Comput.* **13**, 937 (2013).
 - [18] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
 - [19] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, *Rev. Mod. Phys.* **82**, 2313 (2010).
 - [20] H. Weimer, N. Y. Yao, C. R. Laumann, and M. D. Lukin, Long-Range Quantum Gates using Dipolar Crystals, *Phys. Rev. Lett.* **108**, 100501 (2012).

- [21] B. Yan, S. A. Moses, B. Gadway, J. P. Covey, K. R. A. Hazzard, A. M. Rey, D. S. Jin, and J. Ye, Observation of dipolar spin-exchange interactions with lattice-confined polar molecules, *Nature (London)* **501**, 521 (2013).
- [22] N. Y. Yao, L. Jiang, A. V. Gorshkov, P. C. Maurer, G. Giedke, J. I. Cirac, and M. D. Lukin, Scalable architecture for a room temperature solid-state quantum information processor, *Nat. Commun.* **3**, 800 (2012).
- [23] L. Isenhower, M. Saffman, and K. Mølmer, Multibit CNOT quantum gates via Rydberg blockade, *Quant. Info. Proc.* **10**, 755 (2011).
- [24] K. Mølmer, L. Isenhower, and M. Saffman, Efficient Grover search with Rydberg blockade, *J. Phys. B: At. Mol. Opt. Phys.* **44**, 184016 (2011).
- [25] D. Petrosyan, Dipolar exchange induced transparency with Rydberg atoms, *New J. Phys.* **19**, 033001 (2017).
- [26] J. Gulliksen, D. B. R. Dasari, and K. Mølmer, Characterization of how dissipation and dephasing errors accumulate in quantum computers, *EPJ Quantum Technol.* **2**, 4 (2015).
- [27] M. Müller, I. Leshanovsky, H. Weimer, H. P. Büchler, and P. Zoller, Mesoscopic Rydberg Gate Based on Electromagnetically Induced Transparency, *Phys. Rev. Lett.* **102**, 170502 (2009).
- [28] J. T. Young, P. Bienias, R. Belyansky, A. M. Kaufman, and A. V. Gorshkov, Asymmetric Blockade and Multi-Qubit Gates Via Dipole-Dipole Interactions, *Phys. Rev. Lett.* **127**, 120501 (2021).
- [29] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Parallel Implementation of High-Fidelity Multi-qubit Gates with Neutral Atoms, *Phys. Rev. Lett.* **123**, 170503 (2019).
- [30] Z. Eldredge, Z.-X. Gong, J. T. Young, A. H. Moosavian, M. Foss-Feig, and A. V. Gorshkov, Fast Quantum State Transfer and Entanglement Renormalization Using Long-Range Interactions, *Phys. Rev. Lett.* **119**, 170503 (2017).
- [31] A. Y. Guo, M. C. Tran, A. M. Childs, A. V. Gorshkov, and Z.-X. Gong, Signaling and Scrambling with Strongly Long-Range Interactions, *Phys. Rev. A* **102**, 010401(R) (2020).
- [32] M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov, and A. Lucas, Hierarchy of Linear Light Cones with Long-Range Interactions, *Phys. Rev. X* **10**, 031009 (2020).
- [33] T. Kuwahara and K. Saito, Strictly Linear Light Cones in Long-Range Interacting Systems of Arbitrary Dimensions, *Phys. Rev. X* **10**, 031010 (2020).
- [34] We note that in Refs. [54,55], it was shown that the unbounded quantum fanout could be implemented in trapped-ion systems using a global Molmer-Sorensen gate. However, those methods lead to gate durations that scale superlinearly in the system size in order to achieve a certain value of state fidelity and, as such, do not yield an asymptotic speed-up.
- [35] M. C. Tran, A. Y. Guo, A. Deshpande, A. Lucas, and A. V. Gorshkov, Optimal State Transfer and Entanglement Generation in Power-Law Interacting Systems, *Phys. Rev. X* **11**, 031016 (2021).
- [36] D. Coppersmith, IBM Technical Report No. RC19642 (1994), [arXiv:quant-ph/0201067](https://arxiv.org/abs/quant-ph/0201067).
- [37] A. V. Gorshkov, A. M. Rey, A. J. Daley, M. M. Boyd, J. Ye, P. Zoller, and M. D. Lukin, Alkaline-Earth-Metal Atoms as Few-Qubit Quantum Registers, *Phys. Rev. Lett.* **102**, 110503 (2009).
- [38] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.4.L042016> for more details on the implementation of the fanout protocol in a dipolar system as well as a detailed proof of the Theorem.
- [39] R. Cleve and J. Watrous, Fast parallel circuits for the quantum Fourier transform, in *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, FOCS '00 (IEEE Computer Society, New York, 2000), p. 526.
- [40] H. W. Lenstra and C. Pomerance, A rigorous time bound for factoring integers, *J. Amer. Math. Soc.* **5**, 483 (1992).
- [41] T. J. Osborne, Efficient Approximation of the Dynamics of One-Dimensional Quantum Spin Systems, *Phys. Rev. Lett.* **97**, 157202 (2006).
- [42] We note that a similar result could have been achieved using the fact that constant-depth quantum circuits with two fanouts are not strongly classically simulable unless $P = PP$ [56].
- [43] D. Maslov, Linear depth stabilizer and quantum Fourier transformation circuits with no auxiliary qubits in finite neighbor quantum architectures, *Phys. Rev. A* **76**, 052310 (2007).
- [44] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, *Commun. Math. Phys.* **28**, 251 (1972).
- [45] C.-F. Chen and A. Lucas, Finite Speed of Quantum Scrambling with Long Range Interactions, *Phys. Rev. Lett.* **123**, 250605 (2019).
- [46] See Ref. [32] for a more precise definition. We note that our projector $\mathcal{Q}_r$ is a variation on the $\mathbb{Q}_r$ used in Ref. [32]. Specifically, for $D = 1$, they are related via $\mathcal{Q}_r = \sum_{x: |x| \geq r} \mathbb{Q}_x$.
- [47] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, Locality and Digital Quantum Simulation of Power-Law Interactions, *Phys. Rev. X* **9**, 031006 (2019).
- [48] M. B. Hastings and T. Koma, Spectral gap and exponential decay of correlations, *Commun. Math. Phys.* **265**, 781 (2006).
- [49] M. C. Tran, A. Y. Guo, C. L. Baldwin, A. Ehrenberg, A. V. Gorshkov, and A. Lucas, Lieb-Robinson Light Cone for Power-Law Interactions, *Phys. Rev. Lett.* **127**, 160401 (2021).
- [50] T. Kuwahara and K. Saito, Absence of Fast Scrambling in Thermodynamically Stable Long-Range Interacting Systems, *Phys. Rev. Lett.* **126**, 030604 (2021).
- [51] C.-F. Chen and A. Lucas, Optimal frobenius light cone in spin chains with power-law interactions, *Phys. Rev. A* **104**, 062420 (2021).
- [52] J. M. Epstein and K. B. Whaley, Quantum speed limits for quantum-information-processing tasks, *Phys. Rev. A* **95**, 042314 (2017).
- [53] Specifically, for even n , this mapping is defined by $q_i \mapsto q_{2i-1}$ for $i \leq n/2$ and $q_i \mapsto q_{2(n-i+1)}$ for $i > n/2$.
- [54] D. Maslov and Y. Nam, Use of global interactions in efficient quantum circuit constructions, *New J. Phys.* **20**, 033018 (2018).
- [55] Y. Lu, S. Zhang, K. Zhang, W. Chen, Y. Shen, J. Zhang, J.-N. Zhang, and K. Kim, Global entangling gates on arbitrary ion qubits, *Nature (London)* **572**, 363 (2019).
- [56] Y. Takahashi, T. Yamazaki, and K. Tanaka, Hardness of classically simulating quantum circuits with unbounded Toffoli and fan-out gates, *Quantum Inf. Comput.* **14**, 1149 (2014).