

General-Purpose Quantum Circuit Simulator with Projected Entangled-Pair States and the Quantum Supremacy Frontier

Chu Guo,^{1,*} Yong Liu^{①,2,*} Min Xiong,² Shichuan Xue,² Xiang Fu^{①,2} Anqi Huang,² Xiaogang Qiang^{①,2}
 Ping Xu,² Junhua Liu,^{3,4} Shenggen Zheng,⁵ He-Liang Huang,^{1,6,7} Mingtang Deng,²
 Dario Poletti,^{8,†} Wan-Su Bao,^{1,7,‡} and Junjie Wu^{①,2,§}

¹Henan Key Laboratory of Quantum Information and Cryptography, IEU, Zhengzhou 450001, China

²Institute for Quantum Information & State Key Laboratory of High Performance Computing, College of Computer, National University of Defense Technology, Changsha 410073, China

³Information Systems Technology and Design, Singapore University of Technology and Design, 8 Somapah Road, 487372 Singapore

⁴Quantum Intelligence Lab (QI-Lab), Supremacy Future Technologies (SFT), Guangzhou 511340, China

⁵Center for Quantum Computing, Peng Cheng Laboratory, Shenzhen 518055, China

⁶Hefei National Laboratory for Physical Sciences at Microscale and Department of Modern Physics, University of Science and Technology of China, Hefei, Anhui 230026, China

⁷CAS Centre for Excellence and Synergetic Innovation Centre in Quantum Information and Quantum Physics, University of Science and Technology of China, Hefei, Anhui 230026, China

⁸Science and Math Cluster and EPD Pillar, Singapore University of Technology and Design, 8 Somapah Road, 487372 Singapore



(Received 30 July 2019; published 4 November 2019)

Recent advances on quantum computing hardware have pushed quantum computing to the verge of quantum supremacy. Here, we bring together many-body quantum physics and quantum computing by using a method for strongly interacting two-dimensional systems, the projected entangled-pair states, to realize an effective general-purpose simulator of quantum algorithms. The classical computing complexity of this simulator is directly related to the entanglement generation of the underlying quantum circuit rather than the number of qubits or gate operations. We apply our method to study random quantum circuits, which allows us to quantify precisely the memory usage and the time requirements of random quantum circuits. We demonstrate our method by computing one amplitude for a 7×7 lattice of qubits with depth $(1 + 40 + 1)$ on the Tianhe-2 supercomputer.

DOI: 10.1103/PhysRevLett.123.190501

Quantum computers offer the promise of efficiently solving certain problems that are intractable for classical computers, most famously factorizing large numbers [1–3]. With the rapid progress of various quantum systems towards noisy intermediate-scale quantum computing devices [4–11], we are now on the verge of *quantum supremacy* [12], i.e., demonstrating that a quantum computer has the ability to do a computation that no classical computers can tackle, an important milestone in the field of computer science. Various candidates have been suggested to demonstrate quantum supremacy, such as boson sampling [13,14], the instantaneous quantum polynomial protocol [15,16], and random quantum circuits (RQCs) [3,17] which demand less physical resources and are easier to implement compared to, for instance, factorization.

A central aspect for all these near-term supremacy proof-of-principle computations is to produce a quantum state using as fewer number of qubits as well as quantum gate operations as possible, which would nevertheless be highly entangled and hence difficult to obtain and/or characterize by a classical computer, for instance by sampling from it in the computational basis. In the meanwhile, it is important to find effective ways to simulate accurately quantum algorithms on classical computers, which could be used as a benchmarking baseline and to validate near-term quantum devices. In the field of quantum many-body physics, tensor-network states are often used to efficiently represent quantum states with a sizable amount of entanglement [18,19]. The storage required by these tensor-network states is closely related to the amount of entanglement of the quantum state. Recently, matrix product states (which are one-dimensional tensor networks) have been applied to simulate quantum circuits [20]. However, the performance of matrix product states is much less effective if the underlying quantum system is essentially two dimensional. In this work, we present an efficient and generic quantum circuit simulator based on the projected entangled-pair states

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

(PEPS) [19,21–28], a type of tensor-network quantum states representation designed for two-dimensional lattices. Our PEPS-based simulator is a general-purpose quantum circuit simulator for arbitrary quantum circuits: it stores the full quantum state and it can be readily used to compute single amplitudes, observables, and also perform sequences of quantum measurements.

While the quantum circuit simulator we present can tackle generic circuits, in the following we focus on RQCs. They consist of a series of single and two-qubit gates which are applied to different qubits in a particular order. A group of commuting gates, which can be applied simultaneously, constitutes one layer of the circuit, and the more groups of operations that do not commute, the *deeper* the circuit is. More precisely, for the depth of a circuit we will use the notation $(1 + d + 1)$ where the “1”s indicate the Hadamard gates applied to each site at the beginning and at the end of the calculations, while d is the number of noncommuting layers including controlled-Z (CZ) gates and single-qubit gates applied to different sites. RQCs are the standard benchmark for quantum supremacy as put forth by [3]. The general complexity of quantum supremacy experiments is studied in [29]. For RQCs, it was previously shown in [30] that the complexity scales exponentially with $\min(O(dL_h), O(N))$.

RQCs have thus stimulated the search for efficient classical algorithms which would show where exactly the limits of classical simulations are [17,30–38]. State of the art algorithms can be mainly divided into two categories. (i) The state-vector approach which stores the quantum state as a vector and evolves it directly. For example, in [31] a 45-qubit simulation is reported based on this approach. However, this approach is limited by the number of qubits due to the exponential growth of the Hilbert space. (ii) The tensor-based approach, which represents the quantum states as tensors and specifies the input and output states as rank-1 Kronecker projectors. This approach is less sensitive to the number of qubits and has been pursued more actively. For instance, a full amplitude simulation of a 7×7 circuit to depth $(1 + 39 + 1)$ was implemented in 4.2 h on Sunway TaihuLight supercomputer [33], which, however, exploits the weakness in the original design of RQCs in [3]. Recently, it was proposed to trade circuit fidelity for computational efficiency so as to match the fidelity of a given quantum computer [36,37], and practically compute around 1×10^6 amplitudes of a 7×7 circuit to depth $(1 + 40 + 1)$ with 0.5% circuit fidelity in 2.44 h on Summit supercomputer [38]. Our approach differs from the above approaches in that we use PEPS as the data structure to represent the quantum states. Quantum gate operations as well as quantum projections are adapted accordingly to this new data structure.

Quantum circuit simulator Based on PEPS.—In the following, we consider a two-dimensional rectangular lattice of size $L_v \times L_h$, where L_v and L_h are, respectively, the sizes in the vertical and horizontal directions. We use

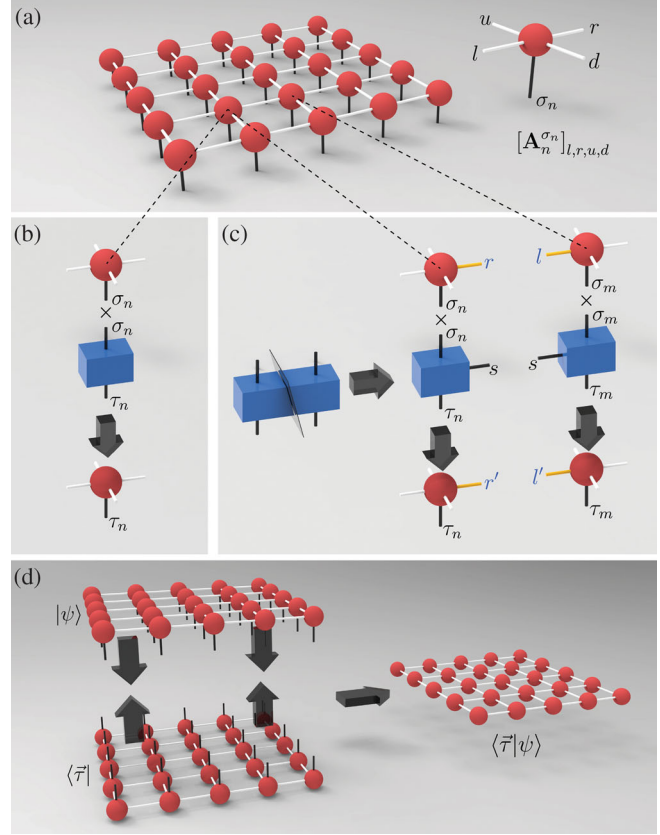


FIG. 1. (a) PEPS on a 5×5 lattice, each qubit of the lattice is represented with a rank-5 tensor $[A_n^{\sigma_n}]_{l,r,u,d}$, where $\sigma_n = 0, 1$ labels the physical dimension and l, r, u, d label the auxiliary dimensions which connect $[A_n^{\sigma_n}]_{l,r,u,d}$ to the tensors on the neighboring sites. (b) Single-qubit gate operation on the PEPS. (c) Two-qubit gate operation on PEPS. (d) Overlapping of two PEPSs by contraction of all the physical dimensions of the two PEPSs and all the auxiliary dimensions inside each PEPS.

$N = L_v L_h$ to denote the total number of qubits. The quantum state on such a lattice can be represented as a PEPS [21,23,24]

$$|\psi\rangle = \sum_{\sigma_1, \dots, \sigma_N} \mathcal{F}(\mathbf{A}_1^{\sigma_1} \mathbf{A}_2^{\sigma_2} \cdots \mathbf{A}_N^{\sigma_N}) |\sigma_1, \sigma_2, \dots, \sigma_N\rangle, \quad (1)$$

where $\mathbf{A}_n^{\sigma_n}$ is a rank-5 tensor with elements $[A_n^{\sigma_n}]_{l,r,u,d}$ at site n , with $\sigma = 0, 1$ corresponding to the physical dimension, and l, r, u, d corresponding to the left, right, up, and down auxiliary dimensions, see Fig. 1(a). The function \mathcal{F} in Eq. (1) indicates the sum over the common auxiliary indices. The bond dimension χ is defined as the maximum size of the four auxiliary dimensions,

$$\chi = \max\{\dim(l), \dim(r), \dim(u), \dim(d)\}, \quad (2)$$

and it characterizes the size of the PEPS.

In the language of PEPS, a single-qubit gate operation $U_{\sigma_n}^r$ on site n only operates locally on the n th tensor $\mathbf{A}_n^{\sigma_n}$ [shown in Fig. 1(b)], which can be written as

$$[\mathbf{A}_n^{\tau_n}]_{l,r,u,d} = \sum_{\sigma_n} U_{\sigma_n}^{\tau_n} [\mathbf{A}_n^{\sigma_n}]_{l,r,u,d}. \quad (3)$$

As we can see from Eq. (3), the size of the local tensor is not affected by a single-qubit gate operation. For a two-qubit gate acting on a horizontally nearest-neighbor pair of qubits (n, m) [shown in Fig. 1(c)], denoted as $O_{\sigma_n, \sigma_m}^{\tau_n, \tau_m}$, we first use a singular value decomposition (SVD) to factorize it into a product of two local tensors

$$\text{SVD}(O_{\sigma_n, \sigma_m}^{\tau_n, \tau_m}) = \sum_s U_{\sigma_n, s}^{\tau_n} V_{s, \sigma_m}^{\tau_m}, \quad (4)$$

where the singular values have been absorbed into U . The size of the auxiliary dimension s is denoted as χ_o , which, for any two-qubit controlled gate, is $\chi_o = 2$. The two local tensors U and V are then applied on the two qubits n and m separately, as single-qubit gate operations

$$[\mathbf{A}_n^{\tau_n}]_{l,r',u,d} = \sum_{\sigma_n} U_{\sigma_n, s}^{\tau_n} [\mathbf{A}_n^{\sigma_n}]_{l,r,u,d}, \quad (5)$$

$$[\mathbf{A}_m^{\tau_m}]_{l',r,u,d} = \sum_{\sigma_m} V_{s, \sigma_m}^{\tau_m} [\mathbf{A}_m^{\sigma_m}]_{l,r,u,d}. \quad (6)$$

Here, we have used the indices $r' = (r, s)$, $l' = (s, l)$, which bundles the two tensor dimensions into one. As a result, χ increases by a factor of χ_o . To keep χ in an affordable size, one would usually use a subsequent singular value decomposition to compress the resulting tensors by throwing away singular values below a suitably chosen threshold. However, we point out that for RQCs we cannot perform such a compression because the distribution of the singular values after the two-qubit gate operation is almost flat, making it impossible for compression (this is also an indication that this problem has large entanglement across the whole circuit). Calculating a single amplitude of the final state $|\psi\rangle$ is done by projecting $|\psi\rangle$ onto a separable PEPS which encodes one spin configuration $|\vec{\tau}\rangle$, and then contracting the resulting tensor network, which can be written as

$$\langle \vec{\tau} | \psi \rangle = \mathcal{F}(\mathbf{E}_1 \mathbf{E}_2 \cdots \mathbf{E}_N), \quad (7)$$

where the rank-4 tensor $[\mathbf{E}_n]_{l,r,u,d} = [\mathbf{A}_n^{\sigma_n = \tau_n}]_{l,r,u,d}$. This calculations are depicted in Fig. 1(d). To this end, we also note that with our method it is also straightforward to simulate sequences of quantum measurements. Concretely, to measure an N th qubit system, we can first compute the probability that a qubit is in state $|0\rangle$ or $|1\rangle$. Then, we use another copy of the wave function (which is stored as PEPS), project the measured qubit in the relevant state, measure another qubit, and so forth. In between different measurements more gates can be applied too, all seamlessly because we can effectively and efficiently compute and store the wave function of the system.

Application to random quantum circuits and complexity analysis.—In the following, we apply our PEPS simulator to study the two-dimensional RQCs of [39,40]. The simulation of this circuit is divided into two parts: (i) circuit evolution and (ii) computing the overlap with randomly selected spin configurations, namely calculating the amplitudes. To quantify the size of the bond dimension required by the tensors, we realize that a single-qubit operation does not affect the size of the tensor it operates on, while a nearest-neighbor two-qubit controlled operation increases the sizes of the two tensors it operates on by a factor of 2 as shown previously [41]. This results in

$$\chi \leq 2^{\lceil d/8 \rceil}, \quad (8)$$

where $\lceil \dots \rceil$ is the ceiling function. The equality in Eq. (8) is reached if the depth d can be divided by 8 (each nearest-neighbor pair of sites will be acted on by a CZ gate in every 8 depths). As can be seen from Eqs. (3), (5), (6), the cost of each gate operation on PEPS scales as $O(\chi^4)$, which is relatively cheap. As a result, circuit evolution can be performed very efficiently. In fact, we can simulate the exact evolution of a 12×12 lattice to a depth $(1 + 40 + 1)$ within minutes on a personal laptop.

In contrast, a well-known result about PEPS is that exactly computing the overlap as in Eq. (7) is an exponentially hard problem [42]. While there exist approximate algorithms to evaluate Eq. (7) which scale polynomially with χ [22,25,26], they are inadequate for RQCs due to the large entanglement of the states produced. In the following, we ignore both the space and time complexity of circuit evolution and only focus on calculating one amplitude, since the cost of the former stage is negligible compared to the latter.

We have developed different strategies to evaluate Eq. (7) efficiently, depending on the shape and size of the lattice. A generic strategy which works for any rectangular lattice has space and time complexities (assuming $L_v \geq L_h$) given by

$$C^s(L_v \times L_h \times d) = 2^{\lceil d/8 \rceil (L_h + 1)}, \quad (9)$$

$$C^t(L_v \times L_h \times d) = (L_h - 2)(L_v - 2)2^{\lceil d/8 \rceil (L_h + 3)}. \quad (10)$$

For square lattices, specialized tensor contraction strategies can be used to further reduce the complexity or for better parallelization (see [40] for details of these strategies). We highlight here that Eqs. (9) and (10) are more accurate estimates for space and time complexities compared to the results of [30], and the exact value will depend on the details of the particular implementation on the hardware. However, these numbers can work as a theoretical approximate benchmarking baseline for achieving quantum supremacy.

To give more precise numbers, using Eqs. (9) and (10), we can evaluate that simulating a 8×8 lattice to a depth

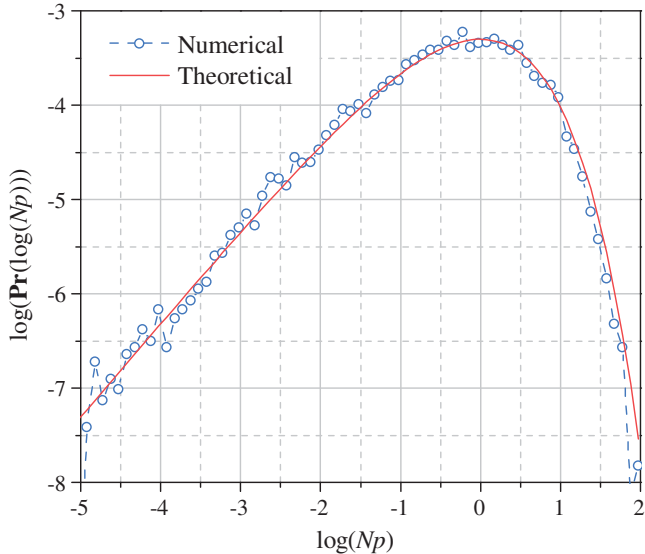


FIG. 2. The blue circles show the log transformed probabilities from calculating 10 000 amplitudes, while the red line is the log transformed Porter-Thomas distribution. The circuit size is 8×8 with a depth $(1 + 25 + 1)$.

$(1 + 40 + 1)$ [same space complexity of a 10×10 circuit to a depth $(1 + 32 + 1)$] would require 32 TB of memory, while simulating a $8 \times l$ (with $l > 8$) lattice to a depth $(1 + 40 + 1)$ would require about 0.5 PB memory. However, simulating a 9×9 lattice with a depth $(1 + 40 + 1)$ would require 16 PB (petabytes) memory and simulating a 12×12 lattice to a depth $(1 + 32 + 1)$ would require 8 PB memory, which are currently out of reach. Our circuit simulator can straightforwardly be extended to other types of two-dimensional lattices including Google Bristlecone QPU architecture. By applying a complexity analysis to this architecture, we find that it only requires less than a manageable 0.6 PB of memory to simulate an RQC with 72 qubits at depth $(1 + 32 + 1)$ (for details of this analysis see [40]).

To demonstrate the performance of our method, we have implemented small scale simulations on a personal computer, which takes less than 1 h to compute one amplitude of a 8×8 circuit to a depth $(1 + 25 + 1)$ for a machine with 2 cores of 2.8 GHz frequency and 16 GB memory. We computed 10 000 amplitudes then plotted the frequency with which each probability of configurations appear. This is represented in Fig. 2 by blue circles, while the red continuous line shows the Porter-Thomas distribution, which is what is expected theoretically.

Our PEPS-based method can be readily scaled up onto a massive parallel computing platform. We implemented the large-scale tensor contractions based on an open-source software package Cyclops Tensor Framework [43]. The massive parallel benchmarking was executed on the Tianhe-2 supercomputer [44]. We have simulated a 7×7 circuit with depth $(1 + 40 + 1)$ and a 10×10 circuit with depth

TABLE I. Large-scale simulation with PEPS-based circuit simulator. The column denoted by “node usage” indicates the number of computing nodes used divided by the total available on Tianhe-2, and the corresponding percentage. “Qubits” and “depth” describe the circuit analyzed while “elapsed time” shows the time required to compute one amplitude.

Node usage	Qubits	Depth	Elapsed time
	7×7	$(1 + 39 + 1)$	9 min
4096/17920, 22%	7×7	$(1 + 40 + 1)$	31 min
	8×8	$(1 + 37 + 1)$	68 min
2048/17920, 11%	9×9	$(1 + 31 + 1)$	22 min
1024/17920, 5%	10×10	$(1 + 26 + 1)$	9 min

$(1 + 26 + 1)$. The simulation of the $7 \times 7 \times (1 + 40 + 1)$ circuit was done on 4096 nodes (22%) of Tianhe-2, taking 31 min and 92.51 TB memory in total [40]. Our large-scale simulation results are listed in Table I.

Conclusions.—In this work we have adapted the projected entangled-pair states representation of quantum states from many-body quantum physics to build a general-purpose quantum circuit simulator. This simulator can be used to store effectively highly entangled wave functions, and it is readily adaptable to compute expectation values or simulate sequential quantum measurements. With this circuit simulator, we have computed an accurate estimate for the space and time complexity analysis of a standard random quantum circuit [39]. Based on this analysis, we point out that simulating an $8 \times l$ circuit to a depth $(1 + 40 + 1)$ or a Bristlecone-72 circuit to a depth $(1 + 32 + 1)$ are within reach of current supercomputing platforms.

We have implemented numerical experiments on a personal computer with a 8×8 circuit to a depth $(1 + 25 + 1)$, and on Tianhe-2 supercomputer with a 10×10 circuit to a depth $(1 + 26 + 1)$, as well as a 7×7 circuit to a depth $(1 + 40 + 1)$. Currently, we compute the amplitudes exactly; however, we could also investigate the trade-off between fidelity and speed, so as to be able to sample many trajectories. For instance, we could reduce the memory requirement of our method by using the “cut” technique in [38], namely mapping a large tensor contraction into summations over many smaller tensor contractions by unraveling several for-loops. More importantly, PEPS-based techniques which are currently used in quantum many-body physics can be transferred to the study of quantum circuits, for example for contractions and the evaluation of expectation values [45]. These investigations, which could be particularly useful for circuits in which the wave function can be effectively compressed, are left for future works, together with the plan to include the effects of noise or errors in order to characterize more closely the actual behavior of a noisy intermediate-scale quantum computer.

We gratefully acknowledge the help from China Greatwall Technology and National Supercomputing

Center in Guangzhou. We thank Sergio Boixo and Giacomo Nannicini for helpful discussions. J. W. acknowledges support from National Natural Science Foundation of China under Grant No. 61632021. C. G. acknowledges support from National Natural Science Foundation of China under Grants No. 11504430 and No. 11805279. P. X. acknowledges support from National Natural Science Foundation of China under Grants No. 11621091 and No. 11690031. H.-L. H. acknowledges support from the Open Research Fund from State Key Laboratory of High Performance Computing of China (Grant No. 201901-01), National Natural Science Foundation of China under Grant No. 11905294, and China Postdoctoral Science Foundation. D. P. acknowledges support from the Singapore Ministry of Education, Singapore Academic Research Fund Tier-II (Project No. MOE2018-T2-2-142).

*These authors contribute equally to this work.

†dario_poletti@sutd.edu.sg

*bws@qiclab.cn

§junjiewu@nudt.edu.cn

- [1] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [2] P. W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE, Piscataway, 1994), pp. 124–134, <https://dx.doi.org/10.1109/SFCS.1994.365700>.
- [3] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [4] A. Lund, M. J. Bremner, and T. Ralph, Quantum sampling problems, BosonSampling and quantum supremacy, *npj Quantum Inf.* **3**, 15 (2017).
- [5] H.-L. Huang, Q. Zhao, X. Ma, C. Liu, Z.-E. Su, X.-L. Wang, L. Li, N.-L. Liu, B. C. Sanders, C.-Y. Lu *et al.*, Experimental Blind Quantum Computing for a Classical Client, *Phys. Rev. Lett.* **119**, 050503 (2017).
- [6] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, *Nature (London)* **551**, 601 (2017).
- [7] H.-L. Huang, X.-L. Wang, P. P. Rohde, Y.-H. Luo, Y.-W. Zhao, C. Liu, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, Demonstration of topological data analysis on a quantum processor, *Optica* **5**, 193 (2018).
- [8] K. Wright, K. Beck, S. Debnath, J. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. Plesch, M. Chmielewski, C. Collins *et al.*, Benchmarking an 11-qubit quantum computer, [arXiv:1903.08181](https://arxiv.org/abs/1903.08181).
- [9] J. Kelly, Z. Chen, B. Chiaro, B. Foxen, J. Martinis, and Q. H. T. Team, Operating and characterizing of a 72 superconducting qubit processor “Bristlecone”: Part 1, in *APS Meeting Abstracts* (2019), Vol. 64, <http://meetings.aps.org/Meeting/MAR19/Session/A42.2>.
- [10] M. Gong, M.-C. Chen, Y. Zheng, S. Wang, C. Zha, H. Deng, Z. Yan, H. Rong, Y. Wu, S. Li *et al.*, Genuine 12-qubit Entanglement on a Superconducting Quantum Processor, *Phys. Rev. Lett.* **122**, 110501 (2019).
- [11] X.-L. Wang, Y.-H. Luo, H.-L. Huang, M.-C. Chen, Z.-E. Su, C. Liu, C. Chen, W. Li, Y.-Q. Fang, X. Jiang *et al.*, 18-Qubit Entanglement with Six Photons’ Three Degrees of Freedom, *Phys. Rev. Lett.* **120**, 260502 (2018).
- [12] J. Preskill, Quantum computing and the entanglement frontier, [arXiv:1203.5813](https://arxiv.org/abs/1203.5813).
- [13] S. Aaronson and A. Arkhipov, The computational complexity of linear optics, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (ACM, New York, 2011), pp. 333–342, <https://dx.doi.org/10.1145/1993636.1993682>.
- [14] J. Wu, Y. Liu, B. Zhang, X. Jin, Y. Wang, H. Wang, and X. Yang, A benchmark test of boson sampling on Tianhe-2 supercomputer, *Natl. Sci. Rev.* **5**, 715 (2018).
- [15] D. Shepherd and M. J. Bremner, Temporally unstructured quantum computation, *Proc. R. Soc. A* **465**, 1413 (2009).
- [16] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy, *Proc. R. Soc. A* **467**, 459 (2011).
- [17] A. Bouldan, B. Fefferman, C. Nirkhe, and U. Vazirani, Quantum supremacy and the complexity of random circuit sampling, [arXiv:1803.04402](https://arxiv.org/abs/1803.04402).
- [18] J. Eisert, Entanglement and tensor network states, [arXiv:1308.3318](https://arxiv.org/abs/1308.3318).
- [19] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys. (Amsterdam)* **349**, 117 (2014).
- [20] A. McCaskey, E. Dumitrescu, M. Chen, D. Lyakh, and T. Humble, Validating quantum-classical programming models with tensor network simulations, *PLoS One* **13**, e0206704 (2018).
- [21] F. Verstraete and J. I. Cirac, Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions, [arXiv:cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066).
- [22] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the Area Law, and the Computational Power of Projected Entangled Pair States, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [23] V. Murg, F. Verstraete, and J. I. Cirac, Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states, *Phys. Rev. A* **75**, 033605 (2007).
- [24] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [25] Z.-C. Gu, M. Levin, and X.-G. Wen, Tensor-entanglement renormalization group approach as a unified method for symmetry breaking and topological phase transitions, *Phys. Rev. B* **78**, 205116 (2008).
- [26] H.-C. Jiang, Z.-Y. Weng, and T. Xiang, Accurate Determination of Tensor Network State of Quantum Lattice Models in Two Dimensions, *Phys. Rev. Lett.* **101**, 090603 (2008).

- [27] Z.-Y. Xie, H.-C. Jiang, Q.N. Chen, Z.-Y. Weng, and T. Xiang, Second Renormalization of Tensor-Network States, *Phys. Rev. Lett.* **103**, 160601 (2009).
- [28] V. Murg, F. Verstraete, and J.I. Cirac, Exploring frustrated spin systems using projected entangled pair states, *Phys. Rev. B* **79**, 195119 (2009).
- [29] S. Aaronson and L. Chen, Complexity-theoretic foundations of quantum supremacy experiments, [arXiv:1612.05903](https://arxiv.org/abs/1612.05903).
- [30] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, and H. Neven, Simulation of low-depth quantum circuits as complex undirected graphical models, [arXiv:1712.05384](https://arxiv.org/abs/1712.05384).
- [31] T. Häner and D. S. Steiger, 0.5 petabyte simulation of a 45-qubit quantum circuit, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (ACM, New York, 2017), p. 33, <https://dx.doi.org/10.1145/3126908.3126947>.
- [32] Z.-Y. Chen, Q. Zhou, C. Xue, X. Yang, G.-C. Guo, and G.-P. Guo, 64-qubit quantum circuit simulation, *Sci. Bull.* **63**, 964 (2018).
- [33] R. Li, B. Wu, M. Ying, X. Sun, and G. Yang, Quantum supremacy circuit simulation on sunway taihulight, [arXiv:1804.04797](https://arxiv.org/abs/1804.04797).
- [34] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, T. Magerlein, E. Solomonik, and R. Wisnieff, Breaking the 49-qubit barrier in the simulation of quantum circuits, [arXiv:1710.05867](https://arxiv.org/abs/1710.05867).
- [35] M.-C. Chen, R. Li, L. Gan, X. Zhu, G. Yang, C.-Y. Lu, and J.-W. Pan, Quantum teleportation-inspired algorithm for sampling large random quantum circuits, [arXiv:1901.05003](https://arxiv.org/abs/1901.05003).
- [36] I.L. Markov, A. Fatima, S. V. Isakov, and S. Boixo, Quantum supremacy is both closer and farther than it appears, [arXiv:1807.10749](https://arxiv.org/abs/1807.10749).
- [37] B. Villalonga, S. Boixo, B. Nelson, C. Henze, E. Rieffel, R. Biswas, and S. Mandrà, A flexible high-performance simulator for the verifying and benchmarking quantum circuits implemented on real hardware, *npj Quantum Inf.* **5**, 86 (2019).
- [38] B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T. S. Humble, R. Biswas, E. G. Rieffel, A. Ho, and S. Mandrà, Establishing the quantum supremacy frontier with a 281 Pflop/s simulation, [arXiv:1905.00444](https://arxiv.org/abs/1905.00444).
- [39] Available on GitHub at <https://github.com/sboixo/GRCS>.
- [40] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.123.190501> for the details on the description of RQCs, the algorithms for computing an amplitude, and the massive parallel benchmarking on Tianhe-2 supercomputer.
- [41] If an iSWAP gate ($|00\rangle\langle 00| + i|01\rangle\langle 10| + i|10\rangle\langle 01| + |11\rangle\langle 11|$) is used instead of a CZ gate, then the size of the two tensors after factorization will increase by a factor 2, effectively halving the depth of the circuit that can be simulated.
- [42] N. Schuch, M.M. Wolf, F. Verstraete, and J.I. Cirac, Computational Complexity of Projected Entangled Pair States, *Phys. Rev. Lett.* **98**, 140506 (2007).
- [43] E. Solomonik, D. Matthews, J. R. Hammond, J. F. Stanton, and J. Demmel, A massively parallel tensor contraction framework for coupled-cluster computations, *J. Parallel Distrib. Comput.* **74**, 3176 (2014).
- [44] X. Liao, L. Xiao, C. Yang, and Y. Lu, MilkyWay-2 supercomputer: system and application, *Front. Comput. Sci.* **8**, 345 (2014).
- [45] M. Lubasch, J.I. Cirac, and M.-C. Bañuls, Unifying projected entangled pair state contractions, *New J. Phys.* **16**, 033014 (2014).