2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)

Quantum advantage with noisy shallow circuits in 3D

Sergey Bravyi
Quantum Computation and Information Group,
IBM T. J. Watson Research Center,
Yorktown Heights, USA

Robert König
Institute for Advanced Study,
Zentrum Mathematik,
Technical University of Munich,
Munich, Germany

David Gosset
Department of Combinatorics and Optimization and
Institute for Quantum Computing,
University of Waterloo,
Waterloo, Canada

Marco Tomamichel
Centre for Quantum Software and Information
and School of Computer Science,
University of Technology Sydney,
Sydney, Australia

***Abstract*—Prior work has shown that there exists a relation problem which can be solved with certainty by a constant-depth quantum circuit composed of geometrically local gates in two dimensions, but cannot be solved with high probability by any classical constant depth circuit composed of bounded fan-in gates. Here we provide two extensions of this result. Firstly, we show that a separation in computational power persists even when the constant-depth quantum circuit is restricted to geometrically local gates in one dimension. The corresponding quantum algorithm is the simplest we know of which achieves a quantum advantage of this type. Our second, main result, is that a separation persists even if the shallow quantum circuit is corrupted by noise. We construct a relation problem which can be solved with near certainty using a noisy constant-depth quantum circuit composed of geometrically local gates in three dimensions, provided the noise rate is below a certain constant threshold value. On the other hand, the problem cannot be solved with high probability by a noise-free classical circuit of constant depth. A key component of the proof is a quantum error-correcting code which admits constant-depth logical Clifford gates and single-shot logical state preparation. We show that the surface code meets these criteria.**

***Keywords*-quantum algorithms; quantum error correction**

I. INTRODUCTION

The appeal of quantum computing lies in the hope that quantum devices may surpass their classical counterparts in certain information processing tasks. Indeed, a universal quantum computer could efficiently solve certain computational problems such as factoring, for which no efficient classical algorithms are known to date. Yet, even an experimental realization of such universal quantum machines – while impressive and potentially useful in applications – would not conclusively establish a computational quantum advantage in the complexity-theoretic sense. Instead, an efficient quantum algorithm must be accompanied with a proof of the classical hardness of the considered problem. For almost any problem of interest, such a proof would itself constitute a major complexity-theoretic advance.

To solidify the theoretical underpinnings of quantum computation, recent work has focused on computational problems where quantum advantage can be established, either conditionally or information-theoretically. Results of the former category rely on certain complexity-theoretic conjectures such as the non-collapse of the polynomial hierarchy as well as specific hardness assumptions for a given problem. For example, so-called IQP circuits and related proposals [1]–[4] provide evidence that classically sampling from the output distribution of certain shallow quantum circuits may be intractable – a key feature first identified by Terhal and DiVincenzo [5] and later strengthened by Aaronson's characterization of postBQP [6]. Some of these works also provide experimental proposals for using a near-term quantum computer to perform a computational task that cannot be performed by any existing classical computer [7]. A rich debate concerning the feasibility of such proposals has prompted improvements to the performance of classical simulation algorithms for quantum computers [8]–[12].

While these results seek to separate efficient (i.e., polynomial-time) quantum computation from efficient classical computation, complementary unconditional results have been obtained for a more narrow question. It has been shown [13] that constant-depth quantum circuits provide a provable computational advantage over constant-depth classical circuits, where both types of circuits are assumed to have bounded fan-in gates. Ref. [13] introduced a computational problem such that

(i) the problem can be solved with certainty by a constant-depth quantum circuit composed of geometrically local gates on a 2D grid of qubits, while
(ii) any classical probabilistic circuit which solves the problem with success probability at least $7/8$ must have depth growing logarithmically with the input size.

This separation also holds in the average-case setting when

2575-8454/19/$31.00 ©2019 IEEE
DOI 10.1109/FOCS.2019.00064

the classical circuit only needs to solve a few instances of the problem that are drawn randomly from a suitable distribution [13, Supplementary Material]. Similar proofs of quantum advantage with associated average-case hardness results for classical circuits have been obtained more recently in [14], [15], see also [16]. In this work we extend these results in two distinct ways.

First, since the quantum algorithm described in Ref. [13] is geometrically local in two dimensions, it is natural to ask whether a provable quantum advantage can also be achieved in a one-dimensional geometry. We answer this question in the affirmative.

Following Ref. [13], below we consider relation problems. Recall that a relation R is defined as a set of valid input-output pairs $(z_{\text{in}}, z_{\text{out}})$, where z_{in} and z_{out} are bit strings of appropriate length. We shall describe a relation by a function $R(z_{\text{in}}, z_{\text{out}})$ that takes values 0 or 1. A classical or quantum circuit is said to solve a relation problem R for some input z_{in} if it outputs a string z_{out} such that $R(z_{\text{in}}, z_{\text{out}}) = 1$. A relation problem has l input-output bits if $|z_{in}| + |z_{out}| = l$.

Result 1 (**Quantum advantage with 1D shallow circuits**). *For each n there exists a relation problem R with roughly n input-output bits and a set of inputs S of size $|S| = \text{poly}(n)$ such that the following holds:*

- *The problem R can be solved with certainty for all inputs by a constant-depth quantum circuit composed of geometrically local gates on a 1D grid.*
- *Any classical probabilistic circuit composed of constant fan-in gates that solves R with probability exceeding 0.9 for a uniformly random input from S must have depth at least $\Omega(\log n)$.*

The formal statement and a proof of this result are given in the full version [17]. As in previous work [13], [14], the separation described in Result 1 is achieved by a quantum algorithm with input/output statistics that are related to those of a certain nonlocal game. Recall that in a nonlocal game, cooperating players are each provided with an input and must each produce an output without communicating with the other players. Their aim is to satisfy a given winning condition, or input/output relation. It is known that quantum players who share entanglement can win certain nonlocal games with higher probability than classical players who share randomness. To prove the above result, we exhibit a constant-depth one-dimensional quantum circuit and a set S of inputs such that the input/output statistics of the circuit given any input in S are directly related to a variant of the well known magic-square game [18], [19]. We further establish that for any classical circuit with low enough depth there are a significant fraction of inputs in S for which the circuit can be viewed as executing a classical strategy for winning this nonlocal game. The result then follows as a result of upper bounds on the winning probability of any classical strategy. The constant-depth quantum circuit which achieves this quantum advantage is a classically controlled Clifford circuit with a particularly simple one-dimensional structure, and may be suitable for a near-term experimental demonstration.

Secondly, we ask if the separation between the power of constant-depth classical and quantum circuits persists even for noisy quantum circuits, i.e., quantum circuits where each qubit/gate can be erroneous with a constant probability. In this paper we compare the computational power of noisy shallow quantum circuits with that of noise-free shallow classical probabilistic circuits. The quantum circuits we consider will be subject to *local stochastic noise* [20]. This noise model assumes that a random Pauli error occurs at each time step in the ideal circuit. The error may affect multiple qubits, but the probability of high-weight errors must be exponentially suppressed. This is quantified by a *noise rate* $p \in [0, 1]$ such that the probability of observing k single-qubit errors at any given subset of k qubits must be at most p^k, see the full version [17] for formal definitions. The (probabilistic) classical circuits we consider will be composed of gates of bounded fan-in.

We note that standard fault-tolerance constructions which emulate a noise-free universal quantum computation using faulty gates and measurements do not directly apply in this setting: these constructions typically lead to non-constant depth circuits. As an example, a quantum error-correcting code with extensive code distance does not have a constant-depth encoding circuit [21]–[23]. Thus, standard quantum error correction methods do not directly provide a generic way to turn a separation such as that established in [13], or the one described in Result 1, into a separation between noisy constant-depth quantum and (noiseless) constant-depth classical circuits. Nevertheless, in this paper we do provide such a generic recipe. Applying the recipe to the separation described in Result 1 we obtain the following.

Result 2 (**Quantum advantage with noisy shallow circuits**). *For each n there exists a relation problem R with roughly n input-output bits and a set of inputs S of size $|S| = poly(n)$ such that the following holds:*

- *The problem R can be solved with probability at least 0.99 for all inputs by a constant-depth quantum circuit composed of geometrically local gates on a 3D grid, subject to local stochastic noise. The noise rate must be below a constant threshold value independent of n.*
- *Any classical probabilistic circuit composed of constant fan-in gates that solves R with probability exceeding 0.9 for a uniformly random input from S must have depth at least*

$$\Omega\left(\frac{\log(n)}{\log(\log(n))}\right).$$

Let us briefly describe the main idea which allows us to convert a quantum advantage with ideal quantum circuits, such as in Result 1, into one with noisy quantum circuits. The recipe is detailed in the full version [17]. It uses the facts that (A) the quantum circuits which achieves the separation are controlled Clifford circuits with a classical control (i.e., for any fixed input a Clifford unitary is applied), and (B) Certain classical computations, such as the decoding needed for quantum error correction, can be incorporated into the definition of the relation problem rather than performed explicitly in the quantum algorithm.

Consider a relation problem R such that a constant-depth controlled-Clifford circuit produces a solution to a given instance with certainty. We are interested in the setting where R cannot be satisfied by any constant-depth classical circuit. Such relations R are provided in Ref. [13] and Result 1. For a fixed input the controlled-Clifford circuit implements a constant-depth Clifford unitary C acting on n qubits followed by measurement of all qubits in the computational basis. Suppose that our goal is to perform a fault-tolerant version of this computation. We imagine encoding each logical qubit using m physical qubits of some CSS-type [24], [25] stabilizer code $\mathcal{Q}_m$.

As noted above, since good codes do not admit constant-depth encoding circuits, we are unable to initialize all logical qubits in the state $|\overline{0}\rangle$. However, we can hope to prepare a version of this state which is corrupted by a known Pauli operator (which may act nontrivially on all physical qubits). To do this we can initialize all m physical qubits, along with a suitable number m_{anc} of ancilla qubits, in the all-zeros state, and then perform a Clifford circuit W which measures all stabilizers of the code to obtain a syndrome s. The resulting state is then

$$(I \otimes |s\rangle\langle s|)\, W|0^m\rangle|0^{m_{\mathrm{anc}}}\rangle \propto \mathsf{Rec}(s)|\overline{0}\rangle|s\rangle. \qquad (1)$$

where the "recovery" Pauli operator $\mathsf{Rec}(s)$ is a function of the syndrome s. We shall be interested in the case when the code $\mathcal{Q}_m$ is a low-density parity-check (LDPC) code, i.e., it has constant weight stabilizer generators such that each qubit is acted upon nontrivially by at most a constant number of them. The syndrome of such codes can be measured by a constant depth Clifford circuit W. Using this procedure we can prepare the desired logical state $|\overline{0}\rangle$ modulo a Pauli recovery operator $\mathsf{Rec}(s)$. The same method can be used to prepare n copies of the state $|\overline{0}\rangle$, modulo a Pauli recovery $\mathsf{Rec}(s)$ acting on nm qubits. Let $\overline{C}$ be the logical version of the Clifford circuit C. Applying this circuit to the prepared logical all-zero state we obtain

$$\overline{C}\mathsf{Rec}(s)|\overline{0}\rangle^{\otimes n} = P(s)\overline{C}|\overline{0}\rangle^{\otimes n} \qquad (2)$$

where $P(s) = \overline{C}\mathsf{Rec}(s)\overline{C}^{\dagger}$ is another Pauli operator which is a simple function of s. Here we require that the logical Clifford $\overline{C}$ is implementable by a constant-depth physical circuit (for example, this holds for any CSS code with transversal logical Hadamard and phase gates). In other words, using such a code $\mathcal{Q}_m$ we are able to implement a logical encoded version of the constant-depth Clifford circuit C, masked by a Pauli operator $P(s)$ that depends on the initial syndrome measurement s obtained in state preparation. The computational basis measurement statistics of the encoded state with the mask Eq. (2) are related to those of the unencoded state with no mask $C|0\rangle^{\otimes n}$ by flipping the bits corresponding to the X-type part of $P(s)$ and then decoding the resulting bit string. Thus we can simulate the desired encoded quantum computation using a constant-depth quantum circuit along with some simple classical postprocessing. If we chose to incorporate this classical postprocessing into the quantum algorithm, it could pose a problem as its depth may not be constant. Happily, it turns out, we can instead modify the definition of the relation problem R to account for the difference.

Now let us consider the noise-tolerance of this procedure. Since the above quantum circuit has a constant depth and uses logical encoded qubits and operations, it can be made to work in the presence of noisy physical gates and measurements, as long as they occur after the state preparation step. Unfortunately, the state preparation step Eq. (1) is not generally fault-tolerant and the whole algorithm can fail due to errors in the measured syndrome s. For example, a single faulty bit of s can potentially damage the recovery operator $\mathsf{Rec}(s)$ at multiple qubits resulting in an uncorrectable error. This can be addressed by using a code $\mathcal{Q}_m$ that admits a so-called *single-shot state preparation* procedure. The latter is closely related to a single-shot error correction [26]. The code $\mathcal{Q}_m$ is said to admit a single-shot state preparation for a single-qubit logical state $\overline{\phi}$ if there exists a number of ancillas m_{anc} (upper bounded by a polynomial function of m) and a constant-depth Clifford circuit W acting on $m + m_{\mathrm{anc}}$ qubits such that, for any local stochastic Pauli error E with noise rate p, we have

$$(I \otimes |s\rangle\langle s|)\, EW|0^m\rangle|0^{m_{\mathrm{anc}}}\rangle \propto F\mathsf{Rec}(s)|\overline{\phi}\rangle|s\rangle.$$

where F is also a local stochastic Pauli error with a possibly larger noise rate $p' \le c_1 p^{c_2}$ for positive constants c_1, c_2. For example, single-shot state basis state preparation allows us to use a constant-depth circuit composed of noisy gates and measurements to prepare a state $F\mathsf{Rec}(s)|\overline{0}\rangle|s\rangle$, where F is a random Pauli error that can be viewed as residual noise. We can also consider single-shot preparation of k-qubit encoded states, with $k > 1$, in which case m should be replaced by mk above.

Putting together these ingredients we obtain a recipe which starts with a relation R defined by the input-output statistics of a constant-depth controlled-Clifford circuit, and converts this "bare relation" into a "noise-tolerant" relation $\mathcal{R}$ that is based on the encoded circuit with single-shot state preparation, and which incorporates the classical postprocessing in its definition. We further show that the input/output statistics of a constant-depth quantum circuit satisfy $\mathcal{R}$, and we show that the depth required for a classical circuit to satisfy $\mathcal{R}$ is comparable to that required to satisfy the bare relation R.

A crucial requirement for the recipe outlined above is the existence of a CSS stabilizer code $\mathcal{Q}_m$ such that elementary logical Clifford gates are implemented by constant-depth Clifford circuits, and which admits a single-shot state preparation procedure. Here we show that the standard surface code satisfies these desiderata. The first requirement follows from previous work [27] which describes how to implement logical single-qubit Hadamard and phase gates in the surface code using constant-depth Clifford circuits. Together with the transversal logical CNOT gate this provides a complete set of Clifford generators which can each be implemented in constant depth. A central technical contribution of our work is to provide a single-shot state preparation procedure for the surface code. Specifically, we show how to prepare a logical Bell state encoded in two identical surface codes.

Result 3 (Single-shot logical Bell state preparation). *For each* $d \geq 4$, *there is a single-shot state preparation procedure for the encoded Bell state* $2^{-1/2}\left(|\overline{00}\rangle + |\overline{11}\rangle\right)$ *shared between two distance-*d *surface codes, each encoding one logical qubit into* $m = d^2 + (d-1)^2$ *physical qubits. The procedure uses a depth-*6 *Clifford circuit* W *composed of geometrically local gates on a 3D grid and computational basis measurements.*

The proof relies crucially on ideas introduced in Ref. [28]. The authors of Ref. [28] showed how to prepare a logical Bell state encoded into a pair of surface codes starting from a 3D grid of qubits initially prepared in a (noisy) cluster state and measuring a suitable subset of qubits. Here we extend the analysis of Ref. [28] and prove that the same protocol yields a single-shot state preparation scheme with a constant error threshold in the presence of local stochastic noise. We leave as an open question whether Result 3 in conjunction with Knill's syndrome measurement method [29], [30] provides a single-shot error correction scheme based on the surface code.

The 3D constant-depth quantum circuit described in Result 2 is obtained by combining the 3D Bell state preparation circuit of Result 3 with the 1D circuit of Result 1 encoded by the surface code (we shall see that the first few gates of this circuit simply prepare Bell states). We show that the encoded 1D circuit can be made geometrically local on a 3D grid using the lattice folding trick introduced in Ref. [27]. The folded encoded 1D circuit uses only nearest-neighbor two-qubit gates on a 3D grid with $O(1)$ qubits per site, as detailed in the full version [17].

ACKNOWLEDGMENT

SB acknowledges support from the IBM Research Frontiers Institute and funding from the MIT-IBM Watson AI Lab under the project Machine Learning in Hilbert space. DG is a CIFAR fellow in the quantum information science program. RK acknowledges support by the Technical University of Munich – Institute of Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under grant agreement no. 291763, by the DFG cluster of excellence 2111 (Munich Center for Quantum Science and Technology), and by the German Federal Ministry of Education through the funding program Photonics Research Germany, contract no. 13N14776 (QCDA-QuantERA). MT would like to thank the Stellenbosch Institute of Advanced Study and the Centre for Quantum Technologies at the National University of Singapore for hosting him while part of this work was completed.

REFERENCES

[1] Michael J Bremner, Ashley Montanaro, and Dan J Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Physical Review Letters*, 117(8):080501, 2016.

[2] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum*, 1:8, April 2017.

[3] Edward Farhi and Aram W Harrow. Quantum supremacy through the quantum approximate optimization algorithm. *arXiv preprint arXiv:1602.07674*, 2016.

[4] Juan Bermejo-Vega, Dominik Hangleiter, Martin Schwarz, Robert Raussendorf, and Jens Eisert. Architectures for quantum simulation showing a quantum speedup. *Phys. Rev. X*, 8:021010, Apr 2018.

[5] Barbara M Terhal and David P DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games. *Quant. Inf. Comp.*, 4(2):134–145, 2004.

[6] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 461, 2063, pages 3473–3482. The Royal Society, 2005.

[7] Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J Bremner, John M Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595, 2018.

[8] Edwin Pednault, John A Gunnels, Giacomo Nannicini, Lior Horesh, Thomas Magerlein, Edgar Solomonik, and Robert Wisnieff. Breaking the 49-qubit barrier in the simulation of quantum circuits. *arXiv preprint arXiv:1710.05867*, 2017.

[9] Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, and Hartmut Neven. Simulation of low-depth quantum circuits as complex undirected graphical models. *arXiv preprint arXiv:1712.05384*, 2017.

[10] Riling Li, Bujiao Wu, Mingsheng Ying, Xiaoming Sun, and Guangwen Yang. Quantum supremacy circuit simulation on Sunway TaihuLight. *arXiv preprint arXiv:1804.04797*, 2018.

[11] Jianxin Chen, Fang Zhang, Mingcheng Chen, Cupjin Huang, Michael Newman, and Yaoyun Shi. Classical simulation of intermediate-size quantum circuits. *arXiv preprint arXiv:1805.01450*, 2018.

[12] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3:181, 2019.

[13] Sergey Bravyi, David Gosset, and Robert Koenig. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, Oct 2018.

[14] Matthew Coudron, Jalex Stark, and Thomas Vidick. Trading locality for time: certifiable randomness from low-depth circuits. *arXiv preprint arXiv:1810.04233*, October 2018.

[15] Franois Le Gall. Average-case quantum advantage with shallow circuits. *arXiv preprint arXiv:1810.12792*, October 2018.

[16] Adam Bene Watts, Robin Kothari, Luke Schaeffer, and Avishay Tal. Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits. Presented at the 22nd Annual Conference on Quantum Information Processing (QIP), January 2019.

[17] Sergey Bravyi, David Gosset, Robert Koenig, and Marco Tomamichel. Quantum advantage with noisy shallow circuits in 3D. *arXiv preprint arXiv:1904.01502*, 2019.

[18] Asher Peres. Incompatible results of quantum measurements. *Physics Letters A*, 151(3-4):107–108, dec 1990.

[19] N. David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical Review Letters*, 65(27):3373–3376, dec 1990.

[20] Omar Fawzi, Antoine Grospellier, and Anthony Leverrier. Constant overhead quantum fault-tolerance with quantum expander codes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 743–754. IEEE, 2018.

[21] Sergey Bravyi, Matthew B. Hastings, and Frank Verstraete. Lieb-Robinson bounds and the generation of correlations and topological quantum order. *Physical Review Letters*, 97(5):050401, 2006.

[22] Lior Eldar and Aram W Harrow. Local Hamiltonians whose ground states are hard to approximate. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 427–438. IEEE, 2017.

[23] Dorit Aharonov and Yonathan Touati. Quantum circuit depth lower bounds for homological codes. *arXiv preprint arXiv:1810.03912*, 2018.

[24] A Robert Calderbank and Peter W Shor. Good quantum error-correcting codes exist. *Physical Review A*, 54(2):1098, 1996.

[25] Andrew Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 452(1954):2551–2577, 1996.

[26] Héctor Bombín. Single-shot fault-tolerant quantum error correction. *Physical Review X*, 5(3):031043, 2015.

[27] Jonathan E Moussa. Transversal Clifford gates on folded surface codes. *Physical Review A*, 94(4):042316, 2016.

[28] Robert Raussendorf, Sergey Bravyi, and Jim Harrington. Long-range quantum entanglement in noisy cluster states. *Phys. Rev. A*, 71:062313, Jun 2005.

[29] E. Knill. Scalable quantum computing in the presence of large detected-error rates. *Phys. Rev. A*, 71:042322, Apr 2005.

[30] Daniel Gottesman and Isaac L. Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402:390–393, 1999.