Random Circuit Sampling

Quynh T. Nguyen[1]
[1]*Department of Physics and Department of EECS, MIT*

(Dated: May 12, 2022)

Random Circuit Sampling (RCS) is a leading proposal for demonstrating quantum computational advantage–the event when a noisy intermediate-scale quantum (NISQ) computer efficiently performs a computational task that is intractable on classical computers. To achieve this goal requires one to carefully design a problem that can be implementable and verifiable on NISQ devices and, furthermore, classically hard with respect to reasonable complexity-theoretic assumptions. We summarize the recent experimental results by Google and USTC, then we overview complexity-theoretic evidence for the classical hardness of RCS.

I. INTRODUCTION

Quantum advantage (quantum supremacy) [1] refers to the event when a quantum computer efficiently performs a computational task which would take a prohibitively long (e.g. exponential) time on any existing classical computers. Due to its profound implications, quantum advantage is considered a milestone in the fields of quantum computation and quantum technology. In terms of the theory of computation, if successfully implemented, it would serve as a refutation of the extended Church-Turing thesis [2], whereby demonstrating the power of quantum computation. In terms of physics and engineering, it would be a landmark of quantum technology in building and manipulating complex quantum systems with high precision. It should be noted, however, that quantum advantage experiments often do not concern about the practicality of the problem being solved.

The leading proposals for quantum advantage are based on sampling problems [3–5], which are experimentally feasible on existing noisy intermediate-scale quantum (NISQ) devices with 50-100 qubits. Furthermore, their classical hardness is supported by complexity-theoretic evidence related to the complexity class **#P**, a generalization of the famous **NP**. Working with this complexity class is thus less restrictive than the other assumptions needed for justifying that, say integer factoring, is classically hard.

Random Circuit Sampling (RCS) is currently the most strongly supported proposal by complexity theory and have been experimentally demonstrated on superconducting quantum computers with up to 60 qubits [6–8]. The problem statement of RCS is mathematically simple. All we need to do is to prepare an (approximately) Haar-random quantum circuit on n qubits acting on an initial trivial state, then we sample from the output state of the circuit. This sampling task is simply native for quantum systems and, furthermore, does not require quantum error correction. The size of the Hilbert space grows exponentially in n. Therefore, a quantum circuit on 50-60 qubits would require classical computers to somehow sample from a distribution of $\sim 10^{18}$ possibilities to simulate the quantum circuit, a prohibitive challenge to existing supercomputers.

There are at least two ways in arguing about the classical hardness of RCS. First, one can estimate the runtime of the existing classical algorithms. This approach is somewhat less standard as classical algorithms get better overtime (e.g. tensor network algorithms that do not store the entire quantum state, algorithms using optimized computer architecture, etc.). The second, more fundamental approach is based on complexity theory, where one hopes to prove the hardness in an asymptotic setting (n goes to infinity). In particular, the goal is to show that classical algorithms cannot simulate RCS in subexponential time, which implies that increasing the number of qubits would eventually leave no chance for classical algorithms.

We describe the 2019 RCS experiment of Google [6] in Section II, then overview the progress in proving the classical hardness of RCS in Section III. Finally, we discuss some open questions in the field in section IV.

II. EXPERIMENTAL DEMONSTRATIONS

Google's experiment [6] was performed on the Sycamore chip with 54 superconducting qubits on a 2D rectangular grid. One of the qubits was not operational, so the experiment actually used 53 qubits. All qubits can be tuned individually and are chosen randomly from the set $\{\sqrt{X}, \sqrt{Y}, \sqrt{W}\}$, where $W = (X + Y)/\sqrt{2}$. Two-qubit entangling gates are implemented between nearest-neighbor qubits. The two-qubit gates are of the form

$$\text{fSim}(\theta, \phi) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -i\sin(\theta) & 0 \\ 0 & -i\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{bmatrix}, \quad (1)$$

which is a CZ gate followed by a rotation about $XX + YY$ (a fraction of iSWAP gate) where the rotation angles are chosen to minimize errors–they are not randomized. It is now worth noting two theoretical justifications for this design choice of Google. First, the work of [9] showed that a family an n-qubit quantum circuit with 2-local gates on a 2D lattice (related but not identical to the Sycamore layout) is approximately 2-design if the depth is $O(\sqrt{n})$. In RCS, we are interested in the output probability distribution, thus a 2-design random distribution is sufficient. The depth of Google's circuit is about 20 which can be considered of the order $\sqrt{n}$. Second, the gate set must include non-Clifford gates, as Clifford circuits are efficiently classically simulated due to Gottesman-Knill theorem [10]. The layout diagram from the original paper is shown in Figure 1.

Google's largest experiment is performed on a 53-qubit circuit with 1,133 one-qubit gates and 430 two-qubit gates for a depth of 20. The average errors for one- gates and two-qubit gates are 0.16% and 0.62%, respectively, while the average readout error is 3.8%. Notably, all these low errors are measured when the components operate *simultaneously.* Using a simple intersection probability argument, the authors estimate a total circuit fidelity of 0.2%. Intuitively (but not rigorously in any sense), this can be interpreted as that 99.8% of the samples come from noise (*uniform* distribution) and 0.2% of them come from the ideal quantum circuit distribution. Therefore, a few millions samples are sufficient to resolve the signal from the target quantum distribution, which is conjectured to be classically hard as discussed in Section III.

The fidelity estimate above turned out to be in excellent agreement experimentally with Google's benchmarking score called linear cross-entropy benchmark (linear XEB). Linear XEB is proposed in [5] as a proxy of fidelity and admits the following formula:

$$F_{XEB} = \sum_{x \sim P_{exp}} 2^n P_{id}(x) - 1, \quad (2)$$

where x is a bitstring sampled in the experiment and $P_{id}(x)$ is the probability of outcome x on the ideal circuit (which needs to be computed classically). If the experiment is perfect and the circuit is Haar-random, then $F_{XEB} = 1$ due to the so-called Porter-Thomas property. On the other hand, if P_{exp} is uniformly random (the experiment is all noise), then $F_{XEB} = 0$. Google's goal is to achieve a nontrivial XEB score of $\frac{1}{poly(n)}$. Linear XEB can be robustly approximated with $poly(n)$ samples as opposed to the naive $exp(n)$ sample complexity for fidelity.

Result: ten random circuit instances are generated. For each instance, three million samples are collected in about 200 seconds and the XEB score is computed. They obtained an average XEB score of $(2.24 \pm 0.21) \times 10^{-3}$. It is worth noting here that Google did not compute XEB directly due to the intractability (on their computing clusters) of simulating their full depth-20 53-qubit circuit. They instead performed a reasonable extrapolation of XEB by simulating simplified versions of the circuit (called "patch" and "elided" circuits) at various number of qubits and depths.

Classical runtime estimate: Google claimed that a runtime estimate of 10,000 years would be required for the Schrödinger-Feynman algortithm [11] to sample from their hardest circuit on the IBM Summit supercomputer. This claim has been shown to be an overestimate, see Section IV for a discussion.

More recently, RCS experiments on 56- and 60-qubit circuits have been conducted by the University of Science and Technology of China (USTC) [7, 8], whereby expanding the gap between classical and quantum runtimes. See Appendix A for a summary of their results.

III. CLASSICAL HARDNESS

We now overview the complexity-theoretic evidence of the classical hardness of RCS. Interestingly, the study of quantum circuits has an intimate connection to the complexity class **#P**. By definition, a **#P**-complete problem is to count the number of satisfying assignments of Boolean formulas (#SAT). We can see that this class is at least as hard as **NP** since **NP** only asks whether there exists a satisfying assignment (SAT).

Fact III.1 (Exact, worst-case hardness)**.** *Computing output probabilities of quantum circuits is* ***#P****-hard.*

Proving this fact is rather trivial (see e.g. [12]). One can encode the solution of an n-variable #SAT problem into an output probability of a poly-size quantum circuit which performs the following steps:

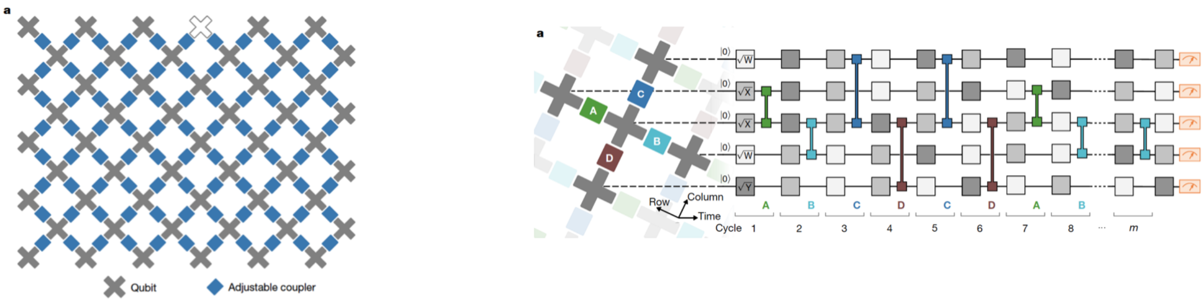

FIG. 1. Layout of the Sycamore chip with 54 superconducting qubits [6]. The nearest-neighbor couplers are controllable two-qubit gates. The two-qubit gates are intertwined by layers of one-qubit gates selected randomly selected from the set $\{\sqrt{X}, \sqrt{Y}, \sqrt{W}\}$, where $W = (X+Y)/\sqrt{2}$.

$$|0^n\rangle\,|0\rangle \to \frac{1}{\sqrt{2^n}}\sum_x |x\rangle\,|0\rangle \;\to \frac{1}{\sqrt{2^n}}\sum_x |x\rangle\,|\phi(x)\rangle\,, \tag{3}$$

where $\phi(x)$ is the SAT formula, i.e. $\phi(x)$ is one if the assignment x is satisfying and zero otherwise. The probability of measuring the second register in state $|1\rangle$ thus encodes the solution of #SAT.

The above fact provides us an intuition why computing *exactly* the probability distribution of a specific (*worst-case*) quantum circuit is hard. However, it does not necessarily state that *sampling* from this distribution is hard. Furthermore, we fundamentally cannot generate a worst-case and noiseless quantum circuit in practice. Indeed, RCS has to do with random (*average-case*) and noisy circuits. Fortunately, Stockmeyer [13] and Aaronson-Arkhipov [4] have shown the following theorem which opens a way to prove an *average-case* hardness of *approximate* sampling.

Theorem III.2 (Sampling-to-computing reduction, informal)**.** *If it is* ***#P****-hard to compute output probabilities of a random quantum circuit up to an additive error of* $O(\varepsilon/2^n)$*, then it is* ***#P****-hard to classically sample from the output distribution of a random circuit with error* ε *in the total variational distance.*

The "if" part of the above theorem remains an unresolved conjecture:

Conjecture III.3 (Approximate, average-case hardness)**.** *It is* ***#P****-hard to compute output probabilities of a random quantum circuit up to an additive error of* $O(\varepsilon/2^n)$*.*

It is good to pause and think about what we have got here and what RCS experiments need to do to demonstrate quantum advantage. If (i) Conjecture III.3 is true and (ii) an RCS experiment can be performed with small error (each outcome's probability has error $\varepsilon/2^n$), then Theorem III.2 says that this experiment is indeed a demonstration of quantum advantage!

We first discuss the progress towards proving Conjecture III.3. The first step is proving approximate, worst-case hardness, which has been shown by a line of work of [3, 4, 12, 13]:

Theorem III.4 (Approximate, worst-case hardness)**.** *Estimating a worst-case circuit output probabilities up to* $O(1)$*-mutiplicative error is* ***#P****-hard, unless the polynomial hierarchy (****PH****) collapses* [1]*.*

The next step is to convert this approximate, worst-case hardness to approximate, average-case hardness *while preserving the error tolerance*. This step is called a worst-to-average-case reduction. Note that Theorem III.4 applies to *multiplicative* error as opposed to the desired *additive* error in Conjecture III.3. This turns out not to be a problem since Haar-random circuits satisfy a property called "anti-concentration", which states that, with high probability over the choice of circuit, $p(x) = \Omega(2^{-n})$ for any possible outcome x. Thus, in the average case $O(1)$-multiplicative error and $O(2^{-n})$-additive error are the same. Unfortunately, we have not been able to achieve this reduction yet. The state-of-art results are due to [14, 15], who constructed a reduction that, unfortunately, suppresses the error tolerance:

[1] **PH** is a hierarchy of generalizations of **NP** and **coNP**. Most of complexity theorists believe in the non-collapse of **PH** (similar to, but not as confident as saying **P**≠**NP**).

Theorem III.5 (Near-exact, average-case hardness)**.** *Estimating random circuit output probabilities up to $O(e^{-m\log m})$-additive error, with m being the number of gates, is* **#P***-hard, unless the polynomial hierarchy collapses.*

In Google's experiment, the circuit depth is $\sqrt{n}$, so $m = n^{3/2}$ and the power factor in this theorem is still polynomially away from the goal of $O(2^{-n})$. Thus, Conjecture III.3 remains unproven.

We now turn to the second requirement for claiming quantum advantage outlined earlier, that is, we need the RCS experiment to have an additive error of only $O(2^{-n})$. Was this achieved by Google? Due to the anti-concentration property, this additive error translates into a *constant* total circuit fidelity. Unfortunately, without correction, the total fidelity of the Sycamore chip is *exponentially small* in the number of gates $O(2^{-m})$. Therefore, Conjecture III.3, if shown to be correct, does not apply to Google's experiment. We refer to the additive error condition in Conjecture III.3 as the low-noise regime, as opposed to the high-noise regime on non-error-corrected quantum circuits. Surprisingly, [14] has shown that Theorem III.5 still holds in the high-noise regime, provided that the noise is gate-independent and below the error detection (not correction!) threshold.

In summary, the classical hardness of RCS remains unproven, but significant progress, especially the work of [14], has provided more evidence and intuition behind Google's computational advantage claim.

IV. DISCUSSION

We have provided an overview of Google's RCS experiment and the complexity-theoretic foundation of RCS. Quantum advantage experiments have been a major step forward, but gaps remain between the theory behind these experiments and their actual implementations. We briefly mention here some open problems/challenges related to RCS not discussed in the main sections. First, a number of classical algorithms have been proposed to challenge Google's 10,000-year claim. The currently best algorithm [16] develops an optimized tensor contraction strategy to simulate Google's experiment in 304 seconds. These algorithms, however, are still essentially exponential time algorithms. Second, the robustness of linear XEB was initially thought to be supported by [11, 17], but the justifications in these works were recently disproved by [18, 19], in which the authors constructed an adversarial noise pattern which leads to an output distribution with high XEB but low fidelity to the ideal circuit.

On a positive note, although this paper implies that Google's RCS experiment does not quite demonstrate an absolute computational advantage yet, one could find an advantage in terms of energy consumption: while Google's experiment draws about 100 kW of power, the Sunway supercomputer used in [16] takes 13 MW. Finally, recent works have proposed applications of RCS in cryptography [20, 21] and Gibbs state preparation [22].

[1] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[2] D. Deutsch, Quantum theory, the church–turing principle and the universal quantum computer, Proc. R. Soc. Lond (1985).

[3] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **467**, 459 (2010).

[4] S. Aaronson and A. Arkhipov, The computational complexity of linear optics, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11 (Association for Computing Machinery, New York, NY, USA, 2011) p. 333–342.

[5] 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, Nature Physics **14**, 595 (2018).

[6] Google, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[7] USTC, Strong quantum computational advantage using a superconducting quantum processor, Physical Review Letters **127**, 10.1103/physrevlett.127.180501 (2021).

[8] USTC, Quantum computational advantage via 60-qubit 24-cycle random circuit sampling (2021).

[9] A. Harrow and S. Mehraban, Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates (2018).

[10] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Physical Review A **70**, 10.1103/physreva.70.052328 (2004).

[11] S. Aaronson and L. Chen, Complexity-theoretic foundations of quantum supremacy experiments (2016).

[12] B. M. Terhal and D. P. DiVincenzo, Adap-

tive quantum computation, constant depth quantum circuits and arthur-merlin games 10.48550/ARXIV.QUANT-PH/0205133 (2002).

[13] L. Stockmeyer, On approximation algorithms for # p, SIAM Journal on Computing **14**, 849 (1985).

[14] A. Bouland, B. Fefferman, Z. Landau, and Y. Liu, Noise and the frontier of quantum supremacy, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (2022) pp. 1308–1317.

[15] Y. Kondo, R. Mori, and R. Movassagh, Quantum supremacy and hardness of estimating output probabilities of quantum circuits, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (2022) pp. 1296–1307.

[16] Y. A. Liu, X. L. Liu, F. N. Li, H. Fu, Y. Yang, J. Song, P. Zhao, Z. Wang, D. Peng, H. Chen, C. Guo, H. Huang, W. Wu, and D. Chen, Closing the "quantum supremacy" gap, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (ACM, 2021).

[17] S. Aaronson and S. Gunn, On the classical hardness of spoofing linear cross-entropy benchmarking (2019).

[18] B. Barak, C.-N. Chou, and X. Gao, Spoofing linear cross-entropy benchmarking in shallow quantum circuits (2020).

[19] X. Gao, M. Kalinowski, C.-N. Chou, M. D. Lukin, B. Barak, and S. Choi, Limitations of linear cross-entropy as a measure for quantum advantage (2021).

[20] Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, A cryptographic test of quantumness and certifiable randomness from a single quantum device (2018).

[21] L. Chen and R. Movassagh, Quantum merkle trees (2021).

[22] O. Shtanko and R. Movassagh, Algorithms for gibbs state preparation on noiseless and noisy random quantum circuits (2021).

Appendix A: USTC's experiments

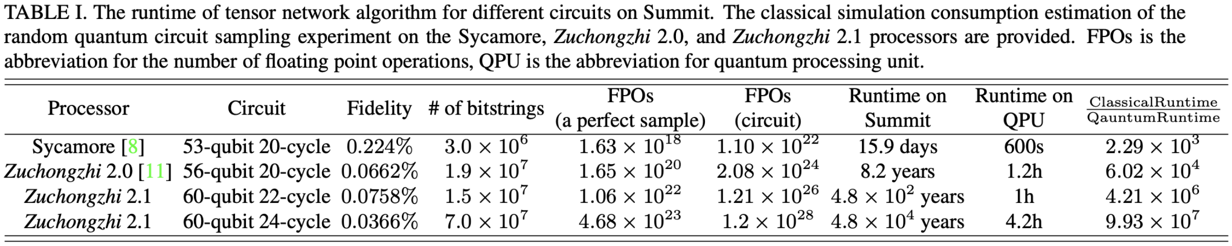

TABLE I. The runtime of tensor network algorithm for different circuits on Summit. The classical simulation consumption estimation of the random quantum circuit sampling experiment on the Sycamore, *Zuchongzhi* 2.0, and *Zuchongzhi* 2.1 processors are provided. FPOs is the abbreviation for the number of floating point operations, QPU is the abbreviation for quantum processing unit.

Processor	Circuit	Fidelity	# of bitstrings	FPOs (a perfect sample)	FPOs (circuit)	Runtime on Summit	Runtime on QPU	$\frac{\text{ClassicalRuntime}}{\text{QauntumRuntime}}$
Sycamore [8]	53-qubit 20-cycle	0.224%	3.0×10^6	1.63×10^{18}	1.10×10^{22}	15.9 days	600s	2.29×10^3
Zuchongzhi 2.0 [11]	56-qubit 20-cycle	0.0662%	1.9×10^7	1.65×10^{20}	2.08×10^{24}	8.2 years	1.2h	6.02×10^4
Zuchongzhi 2.1	60-qubit 22-cycle	0.0758%	1.5×10^7	1.06×10^{22}	1.21×10^{26}	4.8×10^2 years	1h	4.21×10^6
Zuchongzhi 2.1	60-qubit 24-cycle	0.0366%	7.0×10^7	4.68×10^{23}	1.2×10^{28}	4.8×10^4 years	4.2h	9.93×10^7

FIG. 2. Recent superconducting RCS results by Google and USTC, table copied from [8].