THEORY OF COMPUTING, Volume 16 (11), 2020, pp. 1–8
www.theoryofcomputing.org

NOTE

On the Classical Hardness of Spoofing Linear Cross-Entropy Benchmarking

Scott Aaronson* Sam Gunn

Received November 1, 2019; Revised February 11, 2020; Published November 2, 2020

Abstract. Recently, Google announced the first demonstration of quantum computational supremacy with a programmable superconducting processor. Their demonstration is based on collecting samples from the output distribution of a noisy random quantum circuit, then applying a statistical test to those samples called Linear Cross-Entropy Benchmarking (Linear XEB). This raises a theoretical question: How hard is it for a classical computer to spoof the results of the Linear XEB test? In this short note, we adapt an analysis of Aaronson and Chen to prove a conditional hardness result for Linear XEB spoofing. Specifically, we show that the problem is classically hard, assuming that there is no efficient classical algorithm that, given a random n-qubit quantum circuit C, estimates the probability of C outputting a specific output string, say 0^n, with mean squared error even slightly better than that of the trivial estimator that always estimates $1/2^n$. Our result automatically encompasses the case of noisy circuits.

*Supported by a Vannevar Bush Fellowship from the US Department of Defense, a Simons Investigator Award, and the Simons "It from Qubit" collaboration.

ACM Classification: F.1.3, F.1.2

AMS Classification: 81P68, 68Q17

Key words and phrases: quantum supremacy, quantum complexity, sampling complexity

 DOI: 10.4086/toc.2020.v016a011

1 Introduction

Quantum computational supremacy refers to the solution of a well-defined computational task by a programmable quantum computer in significantly less time than is required by the best known algorithms running on existing classical computers, for reasons of asymptotic scaling. It is a prerequisite for useful quantum computation, and is therefore seen as a major milestone in the field. The task of sampling from random quantum circuits (called RCS) is one proposal for achieving quantum supremacy [4, 5, 2]. Unlike other proposals such as Boson Sampling [1] and Commuting Hamiltonians [6], RCS involves a universal quantum computer – one theoretically capable of applying any unitary transformation. Furthermore, RCS currently appears to be the easiest proposal to implement at a large enough scale to demonstrate quantum supremacy.

A research team based at Google has announced a demonstration of quantum computational supremacy, by sampling the output distributions of random quantum circuits [3]. To verify that their circuits were working correctly, they tested their samples using Linear Cross-Entropy Benchmarking (Linear XEB). This test simply checks that the observed samples tend to concentrate on the outputs that have higher probabilities under the ideal distribution for the given quantum circuit. More formally, given samples $z_1, \ldots, z_k \in \{0,1\}^n$, Linear XEB entails checking that $\mathbb{E}_i[P(z_i)]$ is greater than some threshold $b/2^n$, where $P(z)$ is the probability of observing z under the ideal distribution. In the regime of 40-50 qubits, these probabilities can be calculated by a classical supercomputer with enough time.

While there is some support for the conjecture that no classical algorithm can efficiently sample from the output distribution of a random quantum circuit [5], less is known about the hardness of directly spoofing a test like Linear XEB. Results about the hardness of sampling are not quite results about the hardness of spoofing Linear XEB; a device could score well on Linear XEB while being far from correct in total variation distance by, for example, always outputting the items with the k highest probabilities.

Under the assumption that the noise in the device is purely “depolarizing" – that a sample from the circuit was sampled correctly with probability $b-1$ and otherwise sampled uniformly at random – there is stronger evidence that it is difficult to spoof Linear XEB. Namely, if there is a classical algorithm for sampling from a quantum circuit with perfectly depolarizing noise in time T, then with the help of an all-powerful but untrusted prover, one can calculate a good estimate for output probabilities in time $10T/(b-1)$ with high probability over circuits. Together with results of [9], it follows that under the Strong Exponential Time Hypothesis there is a quantum circuit from which one cannot classically sample with depolarizing noise in time $(b-1)2^{(1-o(1))n}$ [3]. We are unaware of any evidence that does not depend on such a strong assumption about the noise.

However, Aaronson and Chen were able to prove the hardness of a different, related verification procedure from a strong *hardness* assumption they called the Quantum Threshold Assumption (QUATH) [2]. Informally, QUATH states that it is impossible for a polynomial-time classical algorithm to guess whether a specific output string like 0^n has greater-than-median probability of being observed as the output of a given n-qubit quantum circuit, with success probability $1/2+\Omega(1/2^n)$. They went on to investigate algorithms for breaking QUATH by estimating the output amplitudes of quantum circuits. For certain classes of circuits, output amplitudes can be efficiently calculated, but in general even efficiently sampling from the output distribution is impossible unless the polynomial hierarchy collapses [1, 6]. Aaronson and Chen found an algorithm for calculating amplitudes of arbitrary circuits that runs in time $d^{O(n)}$ and

$poly(n,d)$ space, where d is the circuit depth. This is now used in some state-of-the-art simulations, but is still too slow and of the wrong form to violate QUATH for larger circuits, as there is no way to trade the accuracy for polynomial-time efficiency.

Here, we formulate a slightly different assumption that we call XQUATH and show that it implies the hardness of spoofing Linear XEB. Like QUATH, the new assumption is quite strong, but makes no reference to sampling. In particular, while we don't know a reduction, refuting XQUATH seems essentially as hard as refuting QUATH. Note that our result says nothing, one way or the other, about the possibility of improvements to algorithms for calculating amplitudes. It just says that there's nothing particular to spoofing Linear XEB that makes it easier than nontrivially estimating amplitudes.

Indeed, since the news of the Google group's success broke, at least four results have potentially improved on the classical simulation efficiency, beyond what Google had considered. First, Gray and Kourtis were able to optimize tensor network contraction methods to obtain a faster classical amplitude estimator, though it is not competitive for calculating millions of amplitudes at once [7]. Second, Pednault et al. argued that, by using secondary storage, the largest existing classical supercomputers should be able to simulate the experiments done at Google in a few days [11]. Third, Napp et al. produced an efficient algorithm for approximately simulating average-case quantum circuits from a certain distribution of *constant* depth circuits, which is impossible to efficiently exactly simulate classically in the worst-case unless the polynomial hierarchy collapses [10]. This algorithm is not efficient for circuits as deep as those used by the Google team. Fourth, Zhou et al. used tensor network algorithms to simulate circuits as large as in the Google experiment, but with different 2-qubit gates that were easier to simulate [12]. Our result provides some explanation for why these improvements had to target the general problem of amplitude estimation, rather than doing anything specific to the problem of spoofing Linear XEB.

2 Preliminaries

Throughout this note we will refer to random quantum circuits. Our results apply to circuits chosen from any distribution $\mathcal{D}$ over circuits on n qubits that is unaffected by appending NOT gates to any subset of the qubits at the end of the circuit.[1] For every such distribution there is a corresponding version of XQUATH. For instance, we could consider a distribution where d alternating layers of random single- and neighboring two-qubit gates are applied to a square lattice of n qubits, similar to the Google experiment. Note that the actual distribution in the Google experiment might have been affected by appending NOT gates, but they could have applied random NOT gates to the end of their circuits classically and achieved the same fidelity. If circuits from $\mathcal{D}$ include randomly-chosen NOT gates in the final layer, then $\mathcal{D}$ obviously satisfies our condition.

Our assumption XQUATH states that no efficient classical algorithm can estimate the probability of such a random circuit C outputting 0^n, with mean squared error even slightly lower than the trivial algorithm that always estimates $1/2^n$.

Definition 2.1 (XQUATH, or Linear Cross-Entropy Quantum Threshold Assumption)**.** There is no polynomial-time classical algorithm that takes as input a quantum circuit $C \leftarrow \mathcal{D}$ and produces an

[1] We will also assume that there is an efficient procedure for converting $C \leftarrow \mathcal{D}$ to a new, identically distributed, C' with NOT gates applied to select outputs on C.

estimate p of $p_0 = \Pr[C \text{ outputs } 0^n]$ such that[2]

$$\mathbb{E}[(p_0 - p)^2] = \mathbb{E}[(p_0 - 2^{-n})^2] - \Omega(2^{-3n})$$

where the expectations are taken over circuits C as well as the algorithm's internal randomness.

The simplest way to attempt to refute XQUATH might be to hope that C is near to a circuit that is classically simulable – e. g., if C contains only near-Clifford gates. However, the fraction of such circuits will decay exponentially with the number of *gates* in the circuit, rather than the number of qubits. Alternatively, one might try k random Feynman paths of the circuit, all of which terminate at 0^n, and take the empirical mean over their contributions to the amplitude. This approach will similarly only yield an improvement in mean squared error over the trivial algorithm that decays exponentially with the number of gates. When the number of gates is much larger than $3n$, as in the Google experiment, it is clear that these approaches cannot violate XQUATH. Even the best existing quantum simulation algorithms do not appear to significantly help in refuting XQUATH for reasonable circuit distributions.

The problem XHOG is to generate outputs of a given quantum circuit that have high expected squared-magnitude amplitudes. These outputs are required to be distinct for reasons that will become clear in the proof of Theorem 1.

Problem 2.2 (XHOG, or Linear Cross-Entropy Heavy Output Generation)**.** Given a circuit C, generate k distinct samples $z_1, \dots, z_k$ such that $\mathbb{E}_i[|\langle z_i| C |0^n\rangle|^2] \geq b/2^n$.

The interesting case is when $b > 1$, and we will generally think of b as a constant. Without fault-tolerance, $b - 1$ will quickly become very small for circuits larger than the experiment can handle. This is a difficulty of applying complexity theory to finite experiments, which fail when the problem instance is too large.

When the depth is large enough, the output probabilities p of almost all circuits are empirically observed to be accurately described by the distribution $2^n \mathrm{e}^{-2^n p}$, although this has only been rigorously proven in some special cases [3, 4, 8]. Under this assumption, for observed outputs z from ideal circuits $C \leftarrow \mathcal{D}$ we have

$$\mathbb{E}[|\langle z| C |0^n\rangle|^2] \approx \int_0^\infty \frac{x}{2^n} x \mathrm{e}^{-x} dx = \frac{2}{2^n}$$

So we expect an ideal circuit to solve XHOG with $b \approx 2$, and a noisy circuit to solve XHOG with b slightly larger than 1. Theorem 1 says that, assuming XQUATH, solving XHOG with $b > 1$ is hard to do classically with many samples and high probability. For completeness, we show in the Appendix that with Google's number of samples and estimated circuit fidelity, they would be expected to solve XHOG with sufficiently high probability.

3 The Reduction

We now provide a reduction from the problem in XQUATH to XHOG. Since we only call the XHOG algorithm once in the reduction, and all other steps are efficient, solving XHOG actually requires as many computational steps as solving the problem in XQUATH, minus $O(k)$.

[2] The reason for the bound being 2^{-3n} will emerge from our analysis.

Theorem 3.1. *Assuming XQUATH, no polynomial-time classical algorithm can solve XHOG with probability* $s > \frac{1}{2} + \frac{1}{2b}$*, and*

$$k \geq \frac{1}{((2s-1)b-1)(b-1)}.$$

With $b = 1+\delta$ *and* $s = \frac{1}{2} + \frac{1}{2b} + \varepsilon$*, the right-hand side is approximately* $1/2\varepsilon\delta$*.*

Proof. Suppose that A is a classical algorithm solving XHOG with the parameters above. Given a quantum circuit $C \leftarrow \mathcal{D}$, first draw a uniformly random $z \in \{0,1\}^n$, and apply NOT gates at the end of C on qubits i where $z_i = 1$ to get a circuit C'. According to our assumption on $\mathcal{D}$, C' is distributed exactly the same as C, even conditioned on a particular z. Also, $\langle 0^n | C | 0^n \rangle = \langle z | C' | 0^n \rangle$, so $\Pr[C \text{ outputs } 0^n] = \Pr[C' \text{ outputs } z]$. Call this probability p_0.

Run A on input C' to get $z_1, \ldots, z_k$ with $\mathbb{E}_i[|\langle z_i | C | 0^n \rangle|^2] \geq b2^{-n}$ (when A succeeds). If $z \in \{z_i\}$, then our algorithm outputs $p = b2^{-n}$; otherwise it outputs $p = 2^{-n}$.

Let $X = (p_0 - 2^{-n})^2 - (p_0 - p)^2$. Then

$$\begin{aligned}
\mathbb{E}[X \mid z \in \{z_i\} \text{ and } A \text{ succeeded}] &= 2 \cdot 2^{-n}(b-1) \cdot \mathbb{E}[p_0 \mid z \in \{z_i\} \text{ and } A \text{ succeeded}] + 2^{-2n}(1-b^2) \\
&\geq 2 \cdot 2^{-n}(b-1)(b2^{-n}) + 2^{-2n}(1-b^2) \\
&= 2^{-2n}(b-1)^2 \\
\mathbb{E}[X \mid z \in \{z_i\} \text{ and } A \text{ failed}] &= 2 \cdot 2^{-n}(b-1) \cdot \mathbb{E}[p_0 \mid z \in \{z_i\} \text{ and } A \text{ failed}] + 2^{-2n}(1-b^2) \\
&\geq -2^{-2n}(b^2-1)
\end{aligned}$$

Since $\mathbb{E}[X \mid z \notin \{z_i\}] = 0$, and since z is uniformly random even conditioned on the output of A and its success or failure,

$$\begin{aligned}
\mathbb{E}[X] &= 2^{-n}ks \cdot \mathbb{E}[X \mid z \in \{z_i\} \text{ and } A \text{ succeeded}] \\
&\quad + 2^{-n}k(1-s) \cdot \mathbb{E}[X \mid z \in \{z_i\} \text{ and } A \text{ failed}] \\
&\geq 2^{-3n}k((2s-1)b-1)(b-1)
\end{aligned}$$

which is $\Omega(2^{-3n})$ as long as $k \geq 1/((2s-1)b-1)(b-1)$. This completes the proof. □

One simple instance of the theorem is to take $s = \frac{3}{4} + \frac{1}{4b}$ and $k = 2(b-1)^{-2}$. Note that even with $s = 1$, we need $k \geq (b-1)^{-2}$ samples for the proof to work.

In fact, if the number of samples k is much smaller than $(b-1)^{-2}$, then even sampling uniformly at random would pass XHOG with non-negligible probability. This can be seen using the Kullback-Leibler (KL) divergence: For a single sample,

$$\mathrm{KL}(\mathrm{e}^{-p}, p\mathrm{e}^{-p}) \approx \int_0^\infty \mathrm{e}^{-p}(b(p-1)-p+2)\log(b(p-1)-p+2)dp$$

It is not hard to calculate that the Taylor expansion of the above around $b = 1$ is $(b-1)^2/2 + O((b-1)^3)$. By additivity, the KL divergence for k samples is approximately $k(b-1)^2/2$. By Pinsker's inequality, the total variation distance is at most

$$\sqrt{k(b-1)^2/4}$$

Therefore, in order to have total variation distance independent of b, one needs $k \approx (b-1)^{-2}$.

Finally, we would like to be confident that one is solving XHOG with sufficiently high probability s for Theorem 1 to apply, without having to perform the experiment enough times to verify this directly. This is easy to address under mild assumptions. Let $Y = |\langle z| C |0^n\rangle|^2$, where z is sampled from our XHOG device which was given $C \leftarrow \mathcal{D}$. Assuming $\mathbb{E}[Y] \geq (2b-1)/2^n$, Chebyshev's inequality shows that

$$\Pr[\bar{Y}_k \leq b/2^n] \leq \Pr[|\bar{Y}_k - \mathbb{E}[Y]| \geq (b-1)/2^n] \leq \frac{(\sigma 2^n)^2}{k^2(b-1)^2}$$

when Y has standard deviation bounded by σ and $\bar{Y}_k$ is the empirical mean of k samples of Y. So, as long as $\sigma = O(2^{-n})$, one only needs $\Omega((b-1)^{-2})$ samples – a condition we already used to prove Theorem 1.

4 Open Problems

We conclude with two open problems related to our reduction.

- Can the classical hardness of spoofing Linear XEB be based on a more secure assumption? Is there a similar assumption to XQUATH that is *equivalent* to the classical hardness of XHOG?
- Is XQUATH true? What is the relationship of XQUATH to QUATH?

Acknowledgements

We thank Umesh Vazirani, Boaz Barak, Daniel Kane, Ryan O'Donnell, and Patrick Rall for helpful discussions on the subject of this note.

References

[1] SCOTT AARONSON AND ALEX ARKHIPOV: The computational complexity of linear optics. *Theory of Computing*, 9(4):143–252, 2013. Preliminary version in STOC'11. [doi:10.4086/toc.2013.v009a004] 2

[2] SCOTT AARONSON AND LIJIE CHEN: Complexity-theoretic foundations of quantum supremacy experiments. In *Proc. 32nd Comput. Complexity Conf. (CCC'17)*, pp. 22:1–22:67, 2017. [doi:10.4230/LIPIcs.CCC.2017.22, arXiv:1612.05903] 2

[3] FRANK ARUTE ET AL.: Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. [doi:10.1038/s41586-019-1666-5] 2, 4

[4] 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):595–600, 2018. [doi:10.1038/s41567-018-0124-x] 2, 4

[5] ADAM BOULAND, BILL FEFFERMAN, CHINMAY NIRKHE, AND UMESH VAZIRANI: On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15(2):159–163, 2019. Preliminary version in ITCS'19. [doi:10.1038/s41567-018-0318-2] 2

[6] MICHAEL BREMNER, RICHARD JOZSA, AND DAN SHEPHERD: Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proc. Royal Soc. London Ser. A*, 467(2126):459–472, 2011. [arXiv:1005.1407] 2

[7] JOHNNIE GRAY AND STEFANOS KOURTIS: Hyper-optimized tensor network contraction, 2020. [arXiv:2002.01935] 3

[8] ARAM HARROW AND SAEED MEHRABAN: Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates. 2018. [arXiv:1809.06957] 4

[9] CUPJIN HUANG, MICHAEL NEWMAN, AND MARIO SZEGEDY: Explicit lower bounds on strong quantum simulation. *CoRR*, 2018. [arXiv:1804.10368] 2

[10] JOHN NAPP, ROLANDO L. LA PLACA, ALEXANDER M. DALZELL, FERNANDO G. S. L. BRANDAO, AND ARAM W. HARROW: Efficient classical simulation of random shallow 2D quantum circuits. 2019. [arXiv:2001.00021] 3

[11] EDWIN PEDNAULT, JOHN A. GUNNELS, GIACOMO NANNICINI, LIOR HORESH, AND ROBERT WISNIEFF: Leveraging secondary storage to simulate deep 54-qubit Sycamore circuits. 2019. [arXiv:1910.09534] 3

[12] YIQING ZHOU, E. MILES STOUDENMIRE, AND XAVIER WAINTAL: What limits the simulation of quantum computers?, 2020. [arXiv:2002.07730] 3

AUTHORS

Scott Aaronson
Professor
The University of Texas at Austin
Austin, Texas, United States
aaronson@cs.utexas.edu
https://www.scottaaronson.com/

Sam Gunn
Ph. D. student
The University of California at Berkeley
Berkeley, California, United States
gunn@cs.berkeley.edu
https://people.eecs.berkeley.edu/~gunn

ABOUT THE AUTHORS

SCOTT AARONSON received his bachelor's degree from Cornell University and his Ph. D. from UC Berkeley. He is known for his blog and for founding the Complexity Zoo. He publishes often in Theory of Computing.

SAM GUNN is a Computer Science Ph. D. student in the Theory Group at UC Berkeley, where he is advised by Umesh Vazirani. This paper was written while he was an undergraduate at UT Austin, where he was fortunate to be advised by Scott Aaronson.