The Quantum Supremacy Tsirelson Inequality

William Kretschmer[*]

1 Introduction

A team based at Google has claimed the first experimental demonstration of quantum computational supremacy on a programmable device [AAB+19]. The experiment involved *random circuit sampling*, where the task is to sample (with nontrivial fidelity) from the output distribution of a quantum circuit containing random 1- and 2-qubit gates. To verify their experiment, they used the so-called *Linear Cross-Entropy Benchmark*, or Linear XEB. Specifically, for an n-qubit quantum circuit C and samples $z_1, \ldots, z_k \in \{0,1\}^n$, the benchmark is given by:

$$b = \frac{2^n}{k} \cdot \sum_{i=1}^{k} |\langle z_i|C|0^n\rangle|^2.$$

The goal is for b to be large with high probability over the choice of the random circuit and the randomness of the sampler, as this demonstrates that the observations tend to concentrate on the outputs that are more likely to be measured under the ideal distribution for C (i.e. the noiseless distribution in which z is measured with probability $|\langle z|C|0^n\rangle|^2$). We formalize this task as the b-XHOG task:

Problem 1 (b-XHOG, or Linear Cross-Entropy Heavy Output Generation)**.** *Given a quantum circuit C on n qubits, output a sample $z \in \{0,1\}^n$ such that $\mathbb{E}\left[|\langle z|C|0^n\rangle|^2\right] \geq \frac{b}{2^n}$, where the expectation is over an implicit distribution over circuits C and over the randomness of the algorithm that outputs z.*

Here, b "large" means b bounded away from 1, as outputting z uniformly at random achieves $b = 1$ on average for any C. On the other hand, if z is drawn from the ideal noiseless distribution for C, and if the random circuits C empirically exhibit the *Porter-Thomas* distribution on output probabilities, then sampling from C achieves $b \approx 2$ [AAB+19, AG20].

Under a strong complexity-theoretic conjecture about the classical hardness of nontrivially estimating output probabilities of quantum circuits, Aaronson and Gunn showed that no classical polynomial-time algorithm can solve b-XHOG for any $b \geq 1 + \frac{1}{\mathrm{poly}(n)}$ on random quantum circuits of polynomial size [AG20]. Thus, a physical quantum computer that solves b-XHOG for $b \geq 1 + \Omega(1)$ is considered strong evidence of quantum computational supremacy.

In this work, we ask: can an efficient quantum algorithm for b-XHOG do substantially better than $b = 2$? That is, what is the largest b for which a polynomial-time quantum algorithm can solve b-XHOG on random circuits?

We refer to our problem as the "quantum supremacy Tsirelson inequality" in reference to the Bell [Bel64] and Tsirelson [CT80] inequalities for quantum nonlocal correlations (for a modern overview, see [CHTW04]). Under this analogy, the quantity b in XHOG plays a similar role as the

[*] University of Texas at Austin. Email: kretsch@cs.utexas.edu.

probability p of winning some nonlocal game. For example, the Bell inequality for the CHSH game [CHSH69] states that no classical strategy can win the game with probability $p > \frac{3}{4}$; we view this as analogous to the conjectured inability of efficient classical algorithms to solve b-XHOG for any $b > 1$. By contrast, a quantum strategy with pre-shared entanglement allows players to win the CHSH game with probability $p = \cos^2\left(\frac{\pi}{8}\right) \approx 0.854 > \frac{3}{4}$. An experiment that wins the CHSH game with probability $p > \frac{3}{4}$, a violation of the Bell inequality, is analogous to an experimental demonstration of b-XHOG for $b > 1$ on a quantum computer that establishes quantum computational supremacy. Finally, the Tsirelson inequality for the CHSH game states that any quantum strategy involving arbitrary pre-shared entanglement wins with probability $p \leq \cos^2\left(\frac{\pi}{8}\right)$. Hence, an upper bound on b for efficient quantum algorithms is the quantum supremacy counterpart to the Tsirelson inequality.

1.1 Our Results

We study the quantum supremacy Tsirelson inequality in the quantum query (or black box) model. That is, we consider distributions over quantum circuits that make queries to a randomized quantum or classical oracle, and ask how many queries to the oracle are needed to solve b-XHOG, in terms of b. Our motivation for studying this problem in the query model is twofold. First, quantum query results often give useful intuition for what to expect in the real world, and can provide insight into why naive algorithmic approaches fail. Second, we view this as an interesting quantum query complexity problem in its own right. Whereas most other quantum query lower bounds involve decision problems [Amb18] or relation problems [Bel15], XHOG is more like a weighted, average-case relation problem, because we only require that $|\langle z|C|0^n\rangle|^2$ be large *on average*. This makes it challenging to apply existing lower bound techniques such as the adversary method, which is only known to tightly characterize decision problems, state conversion problems, and efficiently verifiable relation problems [LMR$^+$11, Bel15].

The XHOG task is well-defined for any distribution of random quantum circuits, so this gives us a choice in selecting the distribution. We focus on three classes of oracle circuits that either resemble random circuits used in practical experiments, or that were previously studied in the context of quantum supremacy. Specifically, we consider (1) canonical state preparation oracles, in which the oracle prepares a Haar-random state $|\psi\rangle$ but otherwise encodes no information about $|\psi\rangle$; (2) Haar-random unitaries; and (3) Fourier Sampling circuits that sample from the Fourier distribution of a random Boolean function.

Our first result is an exponential lower bound on the number of quantum queries needed to solve $(2+\varepsilon)$-XHOG given either of the two types of quantum oracles that we consider:

Theorem 2. *For any $\varepsilon \geq \frac{1}{\mathrm{poly}(n)}$, any quantum query algorithm for $(2+\varepsilon)$-XHOG with query access to either:*

(1) a canonical state preparation oracle for a Haar-random n-qubit state $|\psi\rangle$, or

(2) a Haar-random n-qubit unitary,

requires at least $\Omega\left(\frac{2^{n/4}}{\mathrm{poly}(n)}\right)$ queries.

We do not know if Theorem 2 is optimal, but we show that a simple algorithm based on the quantum collision finding algorithm [BHT97] solves $(2+\Omega(1))$-XHOG using $O\left(2^{n/3}\right)$ queries to either oracle.

Finally, we show that for Fourier Sampling circuits, the naive algorithm of simply running the circuit is optimal among all 1-query algorithms:

Theorem 3. *Any 1-query quantum algorithm for b-XHOG with* Fourier Sampling *circuits achieves $b \leq 3$.*[1]

1.2 Our Techniques

The starting point for our proof of the Tsirelson inequality with a canonical state preparation oracle $\mathcal{O}_\psi$ is a result of Ambainis, Rosmanis, and Unruh [ARU14], which shows that any algorithm that queries $\mathcal{O}_\psi$ can be approximately simulated by a different algorithm that makes no queries, but starts with copies of a resource state that depends on $|\psi\rangle$. This resource state consists of polynomially many (in the number of queries to $\mathcal{O}_\psi$) states of the form $\alpha|\psi\rangle + \beta|\perp\rangle$, i.e. copies of $|\psi\rangle$ in superposition with a canonical orthogonal state $|\perp\rangle$. Our strategy is to show that if any algorithm solves b-XHOG given this resource state, then a similar algorithm solves b-XHOG given copies of $|\psi\rangle$ alone. Then, we prove a lower bound on the number of copies of $|\psi\rangle$ needed to solve b-XHOG. To do so, we argue that if $|\psi\rangle$ is Haar-random, then the best algorithm for b-XHOG given copies of $|\psi\rangle$ is a simple collision-finding algorithm: measure all copies of $|\psi\rangle$ in the computational basis, and output whichever string $z \in \{0,1\}^n$ appears most frequently in the measurement results. For a Haar-random n-qubit state, the chance of seeing any collisions is exponentially unlikely (unless the number of copies of $|\psi\rangle$ is exponentially large in n), and so this does not do much better than measuring a single copy of $|\psi\rangle$ and outputting the result.

To prove the analogous lower bound for b-XHOG with a Haar-random unitary oracle, we show more generally that the canonical state preparation oracles and Haar-random unitary oracles are essentially equivalent as resources, which may be of independent interest. More specifically, we show that for an n-qubit state $|\psi\rangle$, given query access to $\mathcal{O}_\psi$, one can approximately simulate (to exponential precision) a random oracle that prepares $|\psi\rangle$. By "random oracle that prepares $|\psi\rangle$," we mean an n-qubit unitary U_ψ that acts as $U_\psi|0^n\rangle = |\psi\rangle$ but Haar-random everywhere else. We can construct such a U_ψ by taking an arbitrary n-qubit unitary that maps $|0^n\rangle$ to $|\psi\rangle$, then composing it with a Haar-random unitary on the $(2^n - 1)$-dimensional subspace orthogonal to $|0^n\rangle$.

Our lower bound for Fourier Sampling circuits uses an entirely different technique. We use the polynomial method of Beals et al. [BBC+01], which shows that for any quantum algorithm that makes T queries to a classical oracle, the output probabilities of the algorithm can be expressed as degree-$2T$ polynomials in the variables of the classical oracle. Our key observation is that the average linear XEB score achieved by such a quantum query algorithm can *also* be expressed as a polynomial in the variables of the classical oracle. We further observe that this polynomial is constrained by the requirement that the polynomials representing the output probabilities must be nonnegative and sum to 1. This allows us to upper bound the largest linear XEB score achievable by the maximum value of a certain linear program, whose variables are the coefficients of the polynomials that represent the output probabilities of the algorithm. To upper bound this quantity, we exhibit a solution to the dual linear program.

[1] Note that the value of b achieved by the naive quantum algorithm for XHOG depends on the class of circuits used. In contrast to Haar-random circuits that achieve $b \approx 2$, Fourier Sampling circuits achieve $b \approx 3$. This stems from the fact that the amplitudes of a Haar-random quantum state are approximately distributed as *complex* normal random variables, whereas the amplitudes of a state produced by a random Fourier Sampling circuit are approximately distributed as *real* normal random variables.

References

[AAB$^+$19] Frank Arute, Kunal Arya, Ryan Babbush, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. URL: https://doi.org/10.1038/s41586-019-1666-5, doi:10.1038/s41586-019-1666-5. [p. 1]

[AG20] Scott Aaronson and Sam Gunn. On the classical hardness of spoofing linear cross-entropy benchmarking. *Theory of Computing*, 16(11):1–8, 2020. URL: http://www.theoryofcomputing.org/articles/v016a011, doi:10.4086/toc.2020.v016a011. [p. 1]

[Amb18] Andris Ambainis. Understanding quantum algorithms via query complexity. In *Proceedings of the 2018 International Congress of Mathematicians*, volume 3, pages 3249–3270, 2018. [p. 2]

[ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, FOCS 14, page 474483, USA, 2014. IEEE Computer Society. URL: https://doi.org/10.1109/FOCS.2014.57, doi:10.1109/FOCS.2014.57. [p. 3]

[BBC$^+$01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778797, July 2001. URL: https://doi.org/10.1145/502090.502097, doi:10.1145/502090.502097. [p. 3]

[Bel64] John Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, Nov 1964. URL: https://link.aps.org/doi/10.1103/PhysicsPhysiqueFizika.1.195, doi:10.1103/PhysicsPhysiqueFizika.1.195. [p. 1]

[Bel15] Aleksandrs Belovs. Variations on quantum adversary, 2015. arXiv:1504.06943. [p. 2]

[BHT97] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. *SIGACT News*, 28(2):1419, June 1997. URL: https://doi.org/10.1145/261342.261346, doi:10.1145/261342.261346. [p. 2]

[CHSH69] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.*, 23:880–884, Oct 1969. URL: https://link.aps.org/doi/10.1103/PhysRevLett.23.880, doi:10.1103/PhysRevLett.23.880. [p. 2]

[CHTW04] Richard Cleve, Peter Høyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th IEEE Annual Conference on Computational Complexity*, CCC 04, page 236249, USA, 2004. IEEE Computer Society. [p. 1]

[CT80] Boris Cirel'son (Tsirelson). Quantum generalizations of Bell's inequality. *Letters in Mathematical Physics*, 4(2):93–100, 1980. URL: https://doi.org/10.1007/BF00417500, doi:10.1007/BF00417500. [p. 1]

[LMR$^+$11] Troy Lee, Rajat Mittal, Ben W. Reichardt, Robert Špalek, and Mario Szegedy. Quantum query complexity of state conversion. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, FOCS 11, page 344353, USA, 2011. IEEE Computer Society. URL: `https://doi.org/10.1109/FOCS.2011.75`, `doi:10.1109/FOCS.2011.75`. [p. 2]