

The Quantum Supremacy Tsirelson Inequality

William Kretschmer 

University of Texas at Austin, TX, USA

<https://www.cs.utexas.edu/~kretsch/>

kretsch@cs.utexas.edu

Abstract

A leading proposal for verifying near-term quantum supremacy experiments on noisy random quantum circuits is linear cross-entropy benchmarking. For a quantum circuit C on n qubits and a sample $z \in \{0, 1\}^n$, the benchmark involves computing $|\langle z|C|0^n \rangle|^2$, i.e. the probability of measuring z from the output distribution of C on the all zeros input. Under a strong conjecture about the classical hardness of estimating output probabilities of quantum circuits, no polynomial-time classical algorithm given C can output a string z such that $|\langle z|C|0^n \rangle|^2$ is substantially larger than $\frac{1}{2^n}$ (Aaronson and Gunn, 2019). On the other hand, for a random quantum circuit C , sampling z from the output distribution of C achieves $|\langle z|C|0^n \rangle|^2 \approx \frac{2}{2^n}$ on average (Arute et al., 2019).

In analogy with the Tsirelson inequality from quantum nonlocal correlations, we ask: can a polynomial-time quantum algorithm do substantially better than $\frac{2}{2^n}$? We study this question in the query (or black box) model, where the quantum algorithm is given oracle access to C . We show that, for any $\varepsilon \geq \frac{1}{\text{poly}(n)}$, outputting a sample z such that $|\langle z|C|0^n \rangle|^2 \geq \frac{2+\varepsilon}{2^n}$ on average requires at least $\Omega\left(\frac{2^{n/4}}{\text{poly}(n)}\right)$ queries to C , but not more than $O(2^{n/3})$ queries to C , if C is either a Haar-random n -qubit unitary, or a canonical state preparation oracle for a Haar-random n -qubit state. We also show that when C samples from the Fourier distribution of a random Boolean function, the naive algorithm that samples from C is the optimal 1-query algorithm for maximizing $|\langle z|C|0^n \rangle|^2$ on average.

2012 ACM Subject Classification Theory of computation \rightarrow Quantum complexity theory; Theory of computation \rightarrow Oracles and decision trees

Keywords and phrases quantum supremacy, quantum query complexity, random circuit sampling

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.13

Related Version Full version available at <https://arxiv.org/abs/2008.08721>.

Funding *William Kretschmer*: Supported by a Vannevar Bush Fellowship and a National Defense Science and Engineering Graduate (NDSEG) Fellowship from the US Department of Defense.

Acknowledgements Thanks to Scott Aaronson, Sabee Grewal, Sam Gunn, Robin Kothari, Daniel Liang, Patrick Rall, Andrea Rocchetto, and Justin Thaler for helpful discussions and illuminating insights. Thanks also to anonymous reviewers for helpful comments regarding the presentation of this work.

1 Introduction

A team based at Google has claimed the first experimental demonstration of quantum computational supremacy on a programmable device [9]. 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 $z_1, \dots, z_k \in \{0, 1\}^n$, the benchmark is given by:

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



© William Kretschmer;

licensed under Creative Commons License CC-BY

12th Innovations in Theoretical Computer Science Conference (ITCS 2021).

Editor: James R. Lee; Article No. 13; pp. 13:1–13:13

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

13:2 The Quantum Supremacy Tsirelson Inequality

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 $\mathbb{E} [|\langle z|C|0^n\rangle|^2] \geq \frac{b}{2^n}$, 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 $b \approx 2$ [9, 2].

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 $b \geq 1 + \frac{1}{\text{poly}(n)}$ on random quantum circuits of polynomial size [2]. 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? Note that the largest b we could hope for is achieved by the optimal sampler that always outputs the string z maximizing $|\langle z|C|0^n\rangle|^2$. If the random circuits induce a Porter-Thomas distribution on output probabilities, then this solves b -XHOG for $b = \Theta(n)$, because the probabilities of a Porter-Thomas distribution approach i.i.d. exponential random variables (see Fact 10 below). However, finding the largest output probability might be computationally difficult even on a quantum computer, which is why we restrict our attention to *efficient* quantum algorithms.

We refer to our problem as the “quantum supremacy Tsirelson inequality” in reference to the Bell [11] and Tsirelson [18] inequalities for quantum nonlocal correlations (for a modern overview, see [20]). Under this analogy, the quantity b in XHOG plays a similar role as the probability p of winning some nonlocal game. For example, the Bell inequality for the CHSH game [19] 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 $p = \cos^2(\frac{\pi}{8}) \approx 0.854 > \frac{3}{4}$. 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 $p \leq \cos^2(\frac{\pi}{8})$. Hence, an upper bound on b for efficient quantum algorithms is the quantum supremacy counterpart to the Tsirelson inequality. We emphasize that our choice to refer to this as a “Tsirelson inequality” is purely by analogy; we do not claim that the question involving quantum supremacy or the techniques one might use to answer it are otherwise related to quantum nonlocal correlations.

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 [5] or relation problems [12], 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*. Contrast this with the relation problem considered in [1], where the task is to output a z such that $|\langle z|C|0^n\rangle|^2$ is greater than some threshold.

Note that there are known quantum query complexity lower bounds for relation problems [9], and even relation problems where the output is a quantum state [6, 25]. Yet, it is unclear whether existing quantum query lower bound techniques are useful here. Whereas the adversary method tightly characterizes the quantum query complexity of decision problems and state conversion problems [24], it is not even known to characterize the query complexity of relation problems (unless they are efficiently verifiable) [12]. The adversary method appears to be essentially useless for saying anything about XHOG, which is not efficiently verifiable and is not a relation problem in the traditional sense.¹

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.

Canonical State Preparation Oracles

Because the linear cross-entropy benchmark for a circuit C depends only on the state $|\psi\rangle := C|0^n\rangle$ produced by the circuit on the all zeros input, it is natural to consider an oracle \mathcal{O}_ψ that prepares a random state $|\psi\rangle$ without leaking additional information about $|\psi\rangle$. Formally, we choose a Haar-random n -qubit state $|\psi\rangle$, and fix a canonical state $|\perp\rangle$ orthogonal to all n -qubit states.² Then, we take the oracle \mathcal{O}_ψ that acts as $\mathcal{O}_\psi|\perp\rangle = |\psi\rangle$, $\mathcal{O}_\psi|\psi\rangle = |\perp\rangle$, and $\mathcal{O}_\psi|\varphi\rangle = |\varphi\rangle$ for any state $|\varphi\rangle$ that is orthogonal to both $|\perp\rangle$ and $|\psi\rangle$. Equivalently, \mathcal{O}_ψ is the reflection about the state $\frac{|\psi\rangle - |\perp\rangle}{2}$. Finally, we let C be the composition of \mathcal{O}_ψ with any unitary that sends $|0^n\rangle$ to $|\perp\rangle$, so that $C|0^n\rangle = |\psi\rangle$. This model is often chosen when proving lower bounds for quantum algorithms that query state preparation oracles (see e.g. [7, 3, 13]), in part because the ability to simulate \mathcal{O}_ψ follows in a completely black box manner from the ability to prepare $|\psi\rangle$ unitarily without garbage (see Lemma 7 below). Hence, the oracle \mathcal{O}_ψ is “canonical” in the sense that it is uniquely determined by $|\psi\rangle$ and is not any more powerful than any other oracle that prepares $|\psi\rangle$ without garbage.

Haar-Random Unitaries

A random polynomial-size quantum circuit C does not behave like a canonical state preparation oracle: $C|x\rangle$ looks like a random quantum state for *any* computational basis state $|x\rangle$, not just $x = 0^n$. Indeed, random quantum circuits are known to information-theoretically approximate the Haar measure in certain regimes [14, 21], and it seems plausible that they are also computationally difficult to distinguish from the Haar measure. Thus, one could alternatively model random quantum circuits by Haar-random n -qubit unitaries.

¹ As we will see later, however, the polynomial method [10] plays an important role in one of our results.

² We can always assume that a convenient $|\perp\rangle$ exists by extending the Hilbert space, if needed. For example, if $|\psi\rangle$ is an n -qubit state, a natural choice is to encode $|\psi\rangle$ by $|\psi\rangle|1\rangle$ and to choose $|\perp\rangle = |0^n\rangle|0\rangle$.

Fourier Sampling Circuits

Finally, we consider quantum circuits that query a random *classical* oracle. For this, we use FOURIER SAMPLING circuits, which Aaronson and Chen [1] previously studied in the context of proving oracular quantum supremacy for a problem related to XHOG. FOURIER SAMPLING circuits are defined as $H^{\otimes n}U_fH^{\otimes n}$, where U_f is a phase oracle for a uniformly random Boolean function $f : \{0, 1\}^n \rightarrow \{-1, 1\}$. On the all-zeros input, FOURIER SAMPLING circuits output a string $z \in \{0, 1\}^n$ with probability proportional to the squared Fourier coefficient $\hat{f}(z)^2$. This model has the advantage that in principle, one can prove the corresponding quantum supremacy Bell inequality for classical algorithms given query access to f , and that in some cases one can replace f by a pseudorandom function to base quantum supremacy on cryptographic assumptions [1].

Summary of Results

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** (Informal version of Theorem 14 and Theorem 17). *For any $\varepsilon \geq \frac{1}{\text{poly}(n)}$, any quantum query algorithm for $(2 + \varepsilon)$ -XHOG with query access to either:*

- (1) *a canonical state preparation oracle \mathcal{O}_ψ for a Haar-random n -qubit state $|\psi\rangle$, or*
- (2) *a Haar-random n -qubit unitary,*

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

We do not know if Theorem 2 is optimal, but we show in Theorem 15 that a simple algorithm based on the quantum collision finding algorithm [16] solves $(2 + \Omega(1))$ -XHOG using $O(2^{n/3})$ 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** (Informal version of Theorem 19). *Any 1-query quantum algorithm for b -XHOG with FOURIER SAMPLING circuits achieves $b \leq 3$.*³

1.2 Our Techniques

The starting point for our proof of the Tsirelson inequality with a canonical state preparation oracle \mathcal{O}_ψ is a result of Ambainis, Rosmanis, and Unruh [7], which shows that any algorithm that queries \mathcal{O}_ψ 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}_ψ) states of the form $\alpha|\psi\rangle + \beta|\perp\rangle$, i.e. copies of $|\psi\rangle$ in superposition with $|\perp\rangle$. 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

³ 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 $b \approx 3$ (see Proposition 18). 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.

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}_ψ , 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. [10], 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.

Due to space constraints, we defer the proofs to the full version of this paper, available at <https://arxiv.org/abs/2008.08721>.

2 Preliminaries

2.1 Notation

We use $[N]$ to denote the set $\{1, 2, \dots, N\}$. We use $\mathbb{1}$ to denote the identity matrix (of implicit size). We let $\text{TD}(\rho, \sigma)$ denote the trace distance between density matrices ρ and σ , and let $\|A\|_\diamond$ denote the diamond norm of a superoperator A acting on density matrices (see [4] for definitions). For a unitary matrix U , we use $U \cdot U^\dagger$ to denote the superoperator that maps ρ to $U\rho U^\dagger$. In a slight abuse of notation, if A denotes a quantum algorithm (which may consist of unitary gates, measurements, oracle queries, and initialization of ancilla qubits), then we also use A to denote the superoperator corresponding to the action of A on input density matrices.

2.2 Oracles for Quantum States

We frequently consider quantum algorithms that query quantum oracles. In this model, a query to a unitary matrix U consists of a single application of either U , U^\dagger , or controlled versions of U or U^\dagger . We also consider quantum algorithms that make queries to *random* oracles. In analogue with the classical random oracle model, such calls are not randomized

13:6 The Quantum Supremacy Tsirelson Inequality

at each query. Rather, a unitary U is chosen randomly (from some distribution) at the start of the execution of the algorithm, and thereafter all queries for the duration of the algorithm are made to U .

We now define several types of unitary oracles that we will use. These definitions (and associated lemmas giving constructions of them) have appeared implicitly or explicitly in prior work, e.g. [7, 3, 13, 8]. For completeness, we provide proofs of the constructions in the full version.

► **Definition 4.** For an n -qubit quantum state $|\psi\rangle$, the reflection about $|\psi\rangle$, denoted \mathcal{R}_ψ , is the n -qubit unitary $\mathcal{R}_\psi := \mathbb{1} - 2|\psi\rangle\langle\psi|$.

In other words, $|\psi\rangle$ is a -1 eigenstate of \mathcal{R}_ψ , and all states orthogonal to $|\psi\rangle$ are $+1$ eigenstates. Note that some authors define the reflection about $|\psi\rangle$ to be the negation of this operator (e.g. [26, 28, 8]), while others follow our convention (e.g. [15, 23, 3]). This makes little difference, as these definitions are equivalent up to a global phase (or, if using the controlled versions, equivalent up to a Pauli Z gate).

The following lemma shows that \mathcal{R}_ψ can be simulated given any unitary that prepares $|\psi\rangle$ from the all-zeros state, possibly with unentangled garbage.

► **Lemma 5.** Let U be a unitary that acts as $U|0^n\rangle|0^m\rangle = |\psi\rangle|\varphi\rangle$, where $|\psi\rangle$ and $|\varphi\rangle$ are n - and m -qubit states, respectively. Then one can simulate T queries to the reflection \mathcal{R}_ψ using $2T + 1$ queries to U .

► **Definition 6.** For a quantum state $|\psi\rangle$, the canonical state preparation oracle for $|\psi\rangle$, denoted \mathcal{O}_ψ , is the reflection about the state $\frac{|\psi\rangle - |\perp\rangle}{\sqrt{2}}$, where $|\perp\rangle$ is some canonical state orthogonal to $|\psi\rangle$.

Unless otherwise specified, we generally assume that if $|\psi\rangle$ is an n -qubit state, then $|\perp\rangle$ is orthogonal to the space of n -qubit states under a suitable encoding (see Footnote 2).

The next lemma shows that \mathcal{O}_ψ can be simulated from *any* oracle that prepares $|\psi\rangle$ without garbage:

► **Lemma 7.** Let U be an n -qubit unitary that satisfies $U|0^n\rangle = |\psi\rangle$. Then one can simulate T queries to \mathcal{O}_ψ using $4T + 2$ queries to U .

We introduce the notion of a *random* state preparation oracle, which, to our knowledge, is new.

► **Definition 8.** For an n -qubit state $|\psi\rangle$ we define a random state preparation oracle for $|\psi\rangle$, denoted U_ψ , as follows. We fix an arbitrary n -qubit unitary V that satisfies $V|0^n\rangle = |\psi\rangle$, then choose a Haar-random unitary W that acts on the $(2^n - 1)$ -dimensional subspace orthogonal to $|0^n\rangle$ in the space of n -qubit states. Finally, we set $U_\psi = VW$.

The invariance of the Haar measure guarantees that this distribution over U_ψ is independent of the choice of V , and hence this is well-defined. Note that while we often refer to U_ψ as a single unitary matrix, U_ψ really refers to a *distribution* over unitary matrices. Notice also that if $|\psi\rangle$ is distributed as a Haar-random n -qubit state, then U_ψ is distributed as a Haar-random n -qubit unitary.

2.3 Other Useful Facts

We use the following formula for the distance between unitary superoperators in the diamond norm.

► **Fact 9** ([4]). *Let V and W be unitary matrices, and suppose d is the distance between 0 and the polygon in the complex plane whose vertices are the eigenvalues of VW^\dagger . Then*

$$\|V \cdot V^\dagger - W \cdot W^\dagger\|_\diamond = 2\sqrt{1 - d^2}.$$

Finally, we observe that for a Haar-random n -qubit quantum state, the information-theoretically largest linear XEB achievable is $O(n)$.

► **Fact 10**. *Let $|\psi\rangle$ be a Haar-random n -qubit quantum state. Then:*

$$\mathbb{E}_{|\psi\rangle} \left[\max_{z \in \{0,1\}^n} |\langle z|\psi\rangle|^2 \right] \leq \frac{O(n)}{2^n}.$$

3 Canonical State Preparation Oracles

In this section, we prove the quantum supremacy Tsirelson inequality for XHOG with a canonical state preparation oracle for a Haar-random state. We first sketch the important ideas in the proof. At the heart of our proof is the following lemma, due to Ambainis, Rosmanis, and Unruh [7]. It shows that any quantum algorithm that makes queries to a canonical state preparation oracle \mathcal{O}_ψ can be approximately simulated by a quantum algorithm that makes no queries to \mathcal{O}_ψ , and instead receives various copies of $|\psi\rangle$ and superpositions of $|\psi\rangle$ with some canonical orthogonal state.

► **Lemma 11** ([7]). *Let A be a quantum query algorithm that makes T queries to \mathcal{O}_ψ . Then for any k , there is a quantum algorithm B that makes no queries to \mathcal{O}_ψ , and a quantum state $|R\rangle$ of the form:*

$$|R\rangle := \bigotimes_{j=1}^k \alpha_j |\psi\rangle + \beta_j |\perp\rangle$$

such that for any state $|\varphi\rangle$:

$$\text{TD}(A(|\varphi\rangle\langle\varphi|), B(|R\rangle\langle R|, |\varphi\rangle\langle\varphi|)) \leq O\left(\frac{T}{\sqrt{k}}\right).$$

So long as $k \gg T^2$, the output of B will be arbitrarily close to the output of A in trace distance. We will use this and Fact 10 to show that if A solves b -XHOG for some $b > 2$, then so does B . Then, to prove a lower bound on the number of queries T to \mathcal{O}_ψ needed to solve b -XHOG, it suffices to instead lower bound k , the number of states of the form $\alpha_j |\psi\rangle + \beta_j |\perp\rangle$ needed to solve b -XHOG.

When $|\psi\rangle$ is a Haar-random state, notice that the linear XEB depends only on the *magnitude* of the amplitudes in $|\psi\rangle$; the phases are irrelevant. So, when considering algorithms that attempt to solve b -XHOG given only a state $|R\rangle$ of the form used in Lemma 11, we might as well assume that the algorithm randomly reassigns the phases on $|\psi\rangle$. More formally, define the mixed state σ_R as

$$\sigma_R := \mathbb{E}_{\text{diagonal } U} [U^{\otimes k} |R\rangle\langle R| U^{\dagger \otimes k}], \quad (1)$$

where the expectation is over the diagonal unitaries U such that the entries $\langle i|U|i\rangle$ are i.i.d. uniformly random complex phases (and by convention, $\langle \perp|U|\perp\rangle = 1$). Then, the algorithm's average linear XEB score on σ_R is identical to its average linear XEB score on $|R\rangle$, because of the invariance of the Haar measure with respect to phases.

13:8 The Quantum Supremacy Tsirelson Inequality

Next, we observe that one can prepare σ_R by measuring k copies of $|\psi\rangle$ in the computational basis. We prove this in Lemma 12. So, when considering algorithms for XHOG that start with $|R\rangle$, it suffices to instead consider algorithms that simply measure k copies of $|\psi\rangle$ in the computational basis. Such algorithms are much easier to analyze, because once we have measured the k copies of $|\psi\rangle$, we can assume (by convexity) that any optimal such algorithm for XHOG outputs a string z deterministically given the k measurement results. And in that case, clearly the optimal strategy is to output whichever z maximizes the posterior expectation of $|\langle z|\psi\rangle|^2$ given the measurement results. We analyze this strategy in Lemma 13, and show that roughly $2^{n/2}$ copies of $|\psi\rangle$ are needed to solve b -XHOG for b bounded away from 2. The intuition is that the posterior expectation of $|\langle z|\psi\rangle|^2$ increases only when we see z at least twice in the measurement results. However, the probability that any two measurement results are the same is tiny – on the order of 2^{-n} – and so we need to measure at least $2^{n/2}$ copies of $|\psi\rangle$ to see any collisions with decent probability.

We now proceed to proving the necessary lemmas.

► **Lemma 12.** *Let $|\psi\rangle = \sum_{i=1}^N \psi_i|i\rangle$ be an unknown quantum state, and consider a state $|R\rangle$ of the form:*

$$|R\rangle := \bigotimes_{j=1}^k \alpha_j|\psi\rangle + \beta_j|\perp\rangle,$$

where α_j, β_j are known for $j \in [k]$, and the vectors $\{|1\rangle, |2\rangle, \dots, |N\rangle, |\perp\rangle\}$ form an orthonormal basis. Define the mixed state σ_R as above. Then there exists a protocol to prepare σ_R by measuring k copies of $|\psi\rangle$ in the computational basis.

To give some intuition, we note that it is simpler to prove Lemma 12 in the case where $\alpha_j = 1$ for all j . In that case, σ_R can be viewed as an $N^k \times N^k$ density matrix where both the rows and columns are indexed by strings in $[N]^k$. Then, the averaging over diagonal unitaries implies that σ_R is obtained from $(|R\rangle\langle R|)^{\otimes k}$ by zeroing out all entries where the index corresponding to the row is not a reordering of the index corresponding to the column. In fact, one can show that σ_R is expressible as a mixture of pure states, where each pure state is a uniform superposition over basis states that are reorderings of each other. Moreover, the probability associated with each pure state in this mixture is precisely the probability that one of the reorderings is observed when we measure k copies of $|\psi\rangle$ in the computational basis. So, to prepare σ_R , it suffices to measure $|\psi\rangle^{\otimes k}$ and then output the uniform superposition over reorderings of the measurement result.

The proof of Lemma 12 is similar, but we instead have to randomly set some of the measurement results to \perp with probability $|\beta_j|^2$.

Combining Lemma 11 and Lemma 12, we have reduced the problem of lower bounding the number of \mathcal{O}_ψ queries needed to solve b -XHOG, to lower bounding the number of copies of $|\psi\rangle$ needed to solve b -XHOG. The next lemma lower bounds this latter quantity.

► **Lemma 13.** *Let $|\psi\rangle$ be a Haar-random n -qubit quantum state. Consider a quantum algorithm that receives as input $|\psi\rangle^{\otimes k}$ and outputs a string $z \in \{0, 1\}^n$. Then:*

$$\mathbb{E}_{|\psi\rangle, z} [|\langle z|\psi\rangle|^2] \leq \frac{2}{2^n} + \frac{O(k^2)}{4^n}.$$

We note that one should not expect Lemma 13 to be tight for large k (say, $k = \Omega(2^{n/2})$). For example, to achieve $b = 4$, we need at least enough samples to see $m \geq 3$ with good probability. But $\Pr[m \geq 3]$ is negligible unless $k = \Omega(2^{2n/3})$. More generally, a tight bound

on the number of copies of $|\psi\rangle$ needed to achieve a particular value of b seems closely related to the number of measurements of $|\psi\rangle$ needed to see $m \geq b - 1$. This is like a sort of “balls into bins” problem [22, 27] with k balls and 2^n bins, but where the probabilities associated to each bin follow a Dirichlet prior rather than being uniform.

We finally have the tools to prove the main result of this section.

► **Theorem 14.** *Any quantum query algorithm for $(2 + \varepsilon)$ -XHOG with query access to \mathcal{O}_ψ for a Haar-random n -qubit state $|\psi\rangle$ requires $\Omega\left(\frac{2^{n/4}\varepsilon^{5/4}}{n}\right)$ queries.*

Lastly, we give an upper bound on the number of queries needed to nontrivially beat the naive algorithm for XHOG with \mathcal{O}_ψ . In fact, the following algorithm works with *any* oracle that prepares a Haar-random state (including a Haar-random unitary), because the algorithm only needs copies of $|\psi\rangle$ and the ability to perform the reflection \mathcal{R}_ψ . We thank Scott Aaronson for suggesting this approach based on quantum collision-finding.

► **Theorem 15.** *There is a quantum algorithm for $(2 + \Omega(1))$ -XHOG that makes $O(2^{n/3})$ queries to a state preparation oracle for a Haar-random n -qubit state $|\psi\rangle$.*

4 Random State Preparation Oracles

In this section, we show that a canonical state preparation oracle and a random state preparation oracle are essentially equivalent, and use it to prove the quantum supremacy Tsirelson inequality for XHOG with a Haar-random oracle.

By Lemma 7, for a state $|\psi\rangle$, query access to a random state preparation oracle U_ψ implies query access to the canonical state preparation oracle \mathcal{O}_ψ with constant overhead. The reverse direction is less obvious. We know from the definition of U_ψ (Definition 8) that one can simulate U_ψ given *any* n -qubit unitary V that prepares $|\psi\rangle$ from $|0^n\rangle$. So, it is tempting to let $V = \mathcal{O}_\psi$ with $|\perp\rangle = |0^n\rangle$ to argue that \mathcal{O}_ψ allows simulating U_ψ . However, this is only possible if $|0^n\rangle$ is orthogonal to $|\psi\rangle$. And while we previously argued that we can always find a canonical state $|\perp\rangle$ that is orthogonal to $|\psi\rangle$ (Footnote 2), this requires extending the Hilbert space, so that \mathcal{O}_ψ no longer acts on n qubits!

To address this, imagine that we knew an explicit n -qubit state $|\psi^\perp\rangle$ orthogonal to $|\psi\rangle$. Notice that we could perfectly swap $|\psi\rangle$ and $|\psi^\perp\rangle$: the composition $\mathcal{O}_\psi \mathcal{O}_{\psi^\perp} \mathcal{O}_\psi$ sends $|\psi\rangle$ to $|\psi^\perp\rangle$, $|\psi^\perp\rangle$ to $|\psi\rangle$, and acts trivially on all states orthogonal to $|\psi\rangle$ and $|\psi^\perp\rangle$. In particular, this swaps $|\psi\rangle$ and $|\psi^\perp\rangle$ while acting only on the space of n -qubit states. Next, if we know $|\psi^\perp\rangle$ explicitly, we can certainly come up with an n -qubit unitary that sends $|0^n\rangle$ to $|\psi^\perp\rangle$. By composing such a unitary with $\mathcal{O}_\psi \mathcal{O}_{\psi^\perp} \mathcal{O}_\psi$, we are left with an n -qubit unitary that sends $|0^n\rangle$ to $|\psi\rangle$. This is sufficient to construct U_ψ , by Definition 8.

While we do not necessarily have such a state $|\psi^\perp\rangle$, a *random* n -qubit state $|\varphi\rangle$ will be exponentially close to such a $|\psi^\perp\rangle$ with overwhelming probability. The next theorem shows that we can use this observation to *approximately* simulate U_ψ given \mathcal{O}_ψ , by going through the steps above and keeping track of deviation from the ideal construction in terms of $\langle\psi|\varphi\rangle$.

► **Theorem 16.** *Let $|\psi\rangle$ be an n -qubit state. Consider a quantum query algorithm A that makes T queries to U_ψ . Then there is a quantum query algorithm B that makes $2T$ queries to \mathcal{O}_ψ such that:*

$$\left\| \mathbb{E}_{U_\psi} [A] - B \right\| \leq \frac{10T + 4}{2^{n/2}}.$$

The above theorem implies that the oracle \mathcal{O}_ψ in Theorem 14 can be replaced by a Haar-random n -qubit unitary.

13:10 The Quantum Supremacy Tsirelson Inequality

► **Theorem 17.** *Any quantum query algorithm for $(2 + \varepsilon)$ -XHOG with query access to U_ψ for a Haar-random n -qubit state $|\psi\rangle$ (i.e. a Haar-random n -qubit unitary) requires $\Omega\left(\frac{2^{n/4}\varepsilon^{5/4}}{n}\right)$ queries.*

5 Fourier Sampling Circuits

In this section, we prove the quantum supremacy Tsirelson inequality for single-query algorithms over FOURIER SAMPLING circuits.

Throughout this section, we let $N = 2^n$, and let $\mathcal{F}_n := \{f : \{0, 1\}^n \rightarrow \{-1, 1\}\}$ denote the set of all n -bit Boolean functions. Given a function $f \in \mathcal{F}_n$, we define the Fourier coefficient

$$\hat{f}(z) := \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} f(x)(-1)^{x \cdot z}$$

for each $z \in \{0, 1\}^n$. We also define the characters $\chi_z : \{0, 1\}^n \rightarrow \{-1, 1\}$ for each $z \in \{0, 1\}^n$:

$$\chi_z(x) := (-1)^{x \cdot z}.$$

Given oracle access to a function $f \in \mathcal{F}_n$, the FOURIER SAMPLING quantum circuit for f consists of a layer of Hadamard gates, then a single query to f , then another layer of Hadamard gates, so that the resulting circuit samples a string $z \in \{0, 1\}^n$ with probability $\hat{f}(z)^2$. In the context of XHOG, we consider the distribution of FOURIER SAMPLING circuits where the oracle f is chosen uniformly at random from \mathcal{F}_n .

► **Proposition 18.** *FOURIER SAMPLING circuits over n qubits solve $(3 - \frac{2}{2^n})$ -XHOG.*

The following theorem shows the optimality of the 1-query algorithm for XHOG with FOURIER SAMPLING circuits:

► **Theorem 19.** *Any 1-query algorithm for b -XHOG over n -qubit FOURIER SAMPLING circuits satisfies $b \leq 3 - \frac{2}{2^n}$.*

To prove Theorem 19, we use the polynomial method of Beals et al. [10]. Consider a quantum query algorithm that makes T queries to $f \in \mathcal{F}_n$ and outputs a string $z \in \{0, 1\}^n$. The polynomial method implies that for each $z \in \{0, 1\}^n$, the probability that the algorithm outputs z can be expressed as a real multilinear polynomial of degree $2T$ in the bits of f . We write such a polynomial as:

$$p_z(f) = \sum_{S \subset \{0, 1\}^n, |S| \leq 2T} c_{z,S} \cdot \prod_{x \in S} f(x).$$

Then, the expected XEB score of this quantum query algorithm is given by:

$$\frac{1}{2^N} \sum_{f \in \mathcal{F}_n} \sum_{z \in \{0, 1\}^n} p_z(f) \cdot \hat{f}(z)^2. \quad (2)$$

Our key observation is that the quantity (2) is linear in the coefficients $c_{z,S}$. This allows us to express the largest XEB score achievable by polynomials of degree $2T$ as a linear program, with the constraints that the polynomials $\{p_z(f) : z \in \{0, 1\}^n\}$ must represent a probability distribution. Then, the objective value of the linear program can be upper bounded by giving a solution to the dual linear program.

6 Discussion

The most natural question left for future work is whether our bounds could be improved. Our lower bounds for b -XHOG with \mathcal{O}_ψ or U_ψ show that for constant ε , $(2 + \varepsilon)$ -XHOG requires $\Omega\left(\frac{2^{n/4}}{\text{poly}(n)}\right)$ queries to either oracle, while the best upper bound we know of solves $(2 + \varepsilon)$ -XHOG in $O(2^{n/3})$ queries. We conjecture that this upper bound is tight.

One possible approach towards improving the lower bound for b -XHOG with \mathcal{O}_ψ (and by extension, U_ψ) is to use the polynomial method, as we did for the FOURIER SAMPLING lower bound. Indeed, the output probabilities of an algorithm that makes T queries to \mathcal{O}_ψ can be expressed as degree- $2T$ polynomials in the entries of \mathcal{O}_ψ . If we write $|\psi\rangle = \sum_{i=1}^N \alpha_i |i\rangle$, then these are polynomials in the amplitudes $\alpha_1, \dots, \alpha_N$ and the conjugates of the amplitudes $\alpha_1^*, \dots, \alpha_N^*$. Because of the invariance of the Haar measure with respect to phases, and because the linear XEB score depends only on the magnitudes of the amplitudes, we can further assume without loss of generality that the output probabilities are polynomials in the variables $|\alpha_1|^2, \dots, |\alpha_N|^2$, which are equivalently the measurement probabilities of $|\psi\rangle$ in the computational basis. We can also assume that these polynomials are homogeneous, because the input variables satisfy $\sum_{i=1}^N |\alpha_i|^2 = 1$. Like in our FOURIER SAMPLING lower bound, the polynomials are constrained to represent a probability distribution for all valid inputs. However, unlike the FOURIER SAMPLING lower bound, this introduces uncountably many constraints in the primal linear program. It may still be possible to exhibit a solution to the dual linear program if only finitely many of the constraints are relevant (such an approach was used in [17], for example).

Our b -XHOG bound for FOURIER SAMPLING circuits is tight, but it only applies to single-query algorithms. In principle, our lower bound approach via the polynomial method could be generalized to algorithms that make additional queries, by increasing the degree of the polynomials in the linear program and exhibiting another dual solution. The challenge seems to be that the parity constraint on the monomials with nonzero coefficients becomes unwieldy when working with higher degree polynomials.

Beyond possible improvements to the query complexity bounds, it would be interesting to give some evidence that beating the naive XHOG algorithm is hard in the real world. Aaronson and Gunn [2] showed that $(1 + \varepsilon)$ -XHOG is *classically* hard, assuming the classical hardness of nontrivially estimating the output probabilities of random quantum circuits. It is not clear whether a similar argument could work for quantum algorithms, though, because sampling from a random quantum circuit gives a better-than-trivial algorithm for estimating its output probabilities.

References

- 1 Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference, CCC '17*, Dagstuhl, DEU, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- 2 Scott Aaronson and Sam Gunn. On the classical hardness of spoofing linear cross-entropy benchmarking. *Theory of Computing*, 16(11):1–8, 2020. doi:10.4086/toc.2020.v016a011.
- 3 Scott Aaronson, Robin Kothari, William Kretschmer, and Justin Thaler. Quantum Lower Bounds for Approximate Counting via Laurent Polynomials. In Shubhangi Saraf, editor, *35th Computational Complexity Conference (CCC 2020)*, volume 169 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:47, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2020.7.

- 4 Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, STOC '98*, page 20–30, New York, NY, USA, 1998. Association for Computing Machinery. doi:10.1145/276698.276708.
- 5 Andris Ambainis. Understanding quantum algorithms via query complexity. In *Proceedings of the 2018 International Congress of Mathematicians*, volume 3, pages 3249–3270, 2018.
- 6 Andris Ambainis, Loïck Magnin, Martin Roetteler, and Jérémie Roland. Symmetry-assisted adversaries for quantum state generation. In *2011 IEEE 26th Annual Conference on Computational Complexity*, pages 167–177, 2011.
- 7 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 474–483, USA, 2014. IEEE Computer Society. doi:10.1109/FOCS.2014.57.
- 8 Srinivasan Arunachalam, Aleksandrs Belovs, Andrew M. Childs, Robin Kothari, Ansis Rosmanis, and Ronald de Wolf. Quantum Coupon Collector. In Steven T. Flammia, editor, *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, volume 158 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.TQC.2020.10.
- 9 Frank Arute, Kunal Arya, Ryan Babbush, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. doi:10.1038/s41586-019-1666-5.
- 10 Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, July 2001. doi:10.1145/502090.502097.
- 11 John Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, November 1964. doi:10.1103/PhysicsPhysiqueFizika.1.195.
- 12 Aleksandrs Belovs. Variations on quantum adversary, 2015. arXiv:1504.06943.
- 13 Aleksandrs Belovs and Ansis Rosmanis. Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates, 2020. arXiv:2002.06879.
- 14 Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346(2):397–434, 2016. doi:10.1007/s00220-016-2706-8.
- 15 Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In *Quantum Computation and Quantum Information*, volume 305 of *Contemporary Mathematics*, pages 53–74. American Mathematical Society, 2002.
- 16 Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. *SIGACT News*, 28(2):14–19, June 1997. doi:10.1145/261342.261346.
- 17 Mark Bun and Justin Thaler. Dual lower bounds for approximate degree and markov-bernstein inequalities. *Inf. Comput.*, 243(C):2–25, August 2015. doi:10.1016/j.ic.2014.12.003.
- 18 Boris Cirel'son (Tsirelson). Quantum generalizations of Bell's inequality. *Letters in Mathematical Physics*, 4(2):93–100, 1980. doi:10.1007/BF00417500.
- 19 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, October 1969. doi:10.1103/PhysRevLett.23.880.
- 20 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 236–249, USA, 2004. IEEE Computer Society.
- 21 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.
- 22 Norman L. Johnson and Samuel Kotz. *Urn models and their application: an approach to modern discrete probability theory*. Wiley, 1977.

- 23 Shelby Kimmel, Cedric Yen-Yu Lin, Guang Hao Low, Maris Ozols, and Theodore J. Yoder. Hamiltonian simulation with optimal sample complexity. *npj Quantum Information*, 3(1):13, 2017. doi:10.1038/s41534-017-0013-7.
- 24 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 344–353, USA, 2011. IEEE Computer Society. doi:10.1109/FOCS.2011.75.
- 25 Nathan Lindzey and Ansis Rosmanis. A Tight Lower Bound For Non-Coherent Index Erasure. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 59:1–59:37, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ITCS.2020.59.
- 26 Frederic Magniez, Ashwin Nayak, Jeremie Roland, and Miklos Santha. Search via quantum walk. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing, STOC '07*, page 575–584, New York, NY, USA, 2007. Association for Computing Machinery. doi:10.1145/1250790.1250874.
- 27 Martin Raab and Angelika Steger. “Balls into bins” - a simple and tight analysis. In *Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM '98*, pages 159–170, Berlin, Heidelberg, 1998. Springer-Verlag.
- 28 Ben W. Reichardt. Reflections for quantum query algorithms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '11*, page 560–569, USA, 2011. Society for Industrial and Applied Mathematics.