Unconditional Quantum Advantage for Sampling with Shallow Circuits

Adam Bene Watts
University of Calgary, Canada

Natalie Parham
Columbia University, New York, NY, USA

Abstract

Recent work by Bravyi, Gosset, and Koenig showed that there exists a search problem that a constant-depth quantum circuit can solve, but that any constant-depth classical circuit with bounded fan-in cannot. They also pose the question: Can we achieve a similar proof of separation for an input-independent sampling task? In this paper, we show that the answer to this question is yes when the number of random input bits given to the classical circuit is bounded.

We introduce a distribution D_n over $\{0,1\}^n$ and construct a constant-depth uniform quantum circuit family $\{C_n\}_n$ such that C_n samples from a distribution close to D_n in total variation distance. For any $\delta < 1$ we also prove, unconditionally, that any classical circuit with bounded fan-in gates that takes as input $kn + n^\delta$ i.i.d. Bernouli random variables with entropy $1/k$ and produces output close to D_n in total variation distance has depth $\Omega(\log \log n)$. This gives an unconditional proof that constant-depth quantum circuits can sample from distributions that can't be reproduced by constant-depth bounded fan-in classical circuits, even up to additive error. We also show a similar separation between constant-depth quantum circuits with advice and classical circuits with bounded fan-in and fan-out, but access to an unbounded number of i.i.d random inputs.

The distribution D_n and classical circuit lower bounds are inspired by work of Viola, in which he shows a different (but related) distribution cannot be sampled from approximately by constant-depth bounded fan-in classical circuits.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory; Theory of computation → Circuit complexity

Keywords and phrases Circuit Complexity, Sampling Separation, Shallow Quantum Circuits, Unconditional Separations, Complexity of Distributions

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.17

Related Version *Full Version*: `https://arxiv.org/abs/2301.00995`

Acknowledgements The authors would like to thank David Gosset for helpful discussions, and Ansis Rosmanis for sharing an insightful note. They thank Angus Lowe for insightful discussions which motivating the study of classical circuits with biased input in Appendix C of the full version of this paper. They also thank Michael Oliveira for insights regarding compilation of the $U_{m,\theta}$ unitary, which are discussed in Appendix A of the full version of the paper.

1 Introduction

At the heart of quantum information theory lies the remarkable observation that quantum devices can process information in ways that classical devices cannot. This is illustrated strikingly by the work of Bell, which showed that measurements made on spatially separated parts of a quantum system could produce non-classical correlations. More recently, much of the excitement surrounding quantum computers comes from the belief that there are problems, for example factoring [20], which can be solved by quantum computers in polynomial time but which cannot be solved efficiently by classical computers.

While Bell's predictions have been verified experimentally [12], there are significant challenges to demonstrating quantum advantage for complex computational problems such as factoring. The current best known quantum algorithms for these problems require construction of a full scale fault-tolerant quantum computer, which is well beyond our current experimental capabilities. Additionally, while it is widely believed that there is no efficient classical algorithm for factoring, this hardness has not been proven formally. Indeed, proving any problem can be solved efficiently by a quantum computer but not by a classical computer would also require proving $\mathsf{P} \neq \mathsf{PSPACE}$ [25], constituting a major breakthrough in complexity theory.

One approach to demonstrating quantum advantage while avoiding these problems is to study the task of sampling from the output distribution of quantum circuits. In this setting it seems possible that even shallow quantum circuits (that is quantum, circuits whose depth is much less than their length) can perform tasks which are still hard classically. In 2004, Terhal and DiVincenzo provided evidence, later strengthened by Aaronson [1], that there is no polynomial time classical algorithm that takes as input a description of a depth-3 quantum circuit and produces samples from the output distribution of that circuit [22]. More recently, a series of works [7, 3, 6] studied the complexity of sampling from the output distribution of a randomly generated shallow quantum circuit (again given a description of the circuit as input) and gave evidence this task couldn't be performed by classical computers in polynomial time. We refer the reader to [14] for a more complete discussion of these issues.

While these examples are striking, they do have some limitations. As in the case of factoring, the proofs of classical hardness in the results discussed above are conditional, relying on (natural) complexity-theoretic conjectures. More subtly, the presence of noise in real-world experiments means that quantum computers will not sample from the ideal output distribution of quantum circuits *exactly*. Near-term (NISQ [18]) devices will likely only sample from the output distribution of the idealized quantum circuits up to (likely large) *additive error*. Strengthening hardness-of-sampling results of the form described above to this more real-world scenario requires more tenuous complexity-theoretic conjectures.

An alternate approach, pioneered by Bravyi, Gosset, and Konig in [9] is to compare the computational power of shallow quantum circuits to the computational power of similarly restricted classical circuits. This allowed for an *unconditional* separation: in [9] they showed that constant-depth quantum (QNC^0) circuits could solve a relational (search) problem – mapping inputs to valid outputs – in a way that constant-depth, bounded fan-in, classical (NC^0) circuits could not. Later work [26, 13] improved on their result to give separations between QNC^0 circuits and more powerful classes of constant-depth classical circuits, or between quantum circuits and classical circuits even in the presence of noise [8].

One notable feature of all of the separations discussed so far is that they are *input-dependent*. That is, they are based on the classical hardness of mapping some input to some output, e.g., a positive integer to its prime factors in the case of factoring, a circuit description to a sample from its output distribution in the case of circuit sampling problems, or a measurement setting to correlated measurement outcomes in the case of the Bell test. For each of these problems it may be easy to produce a valid output for any fixed input – the classical hardness is in finding a classical process that takes *all* valid inputs to valid outputs.

Many important questions in quantum computing, however, concern operations that are *inputless*. A major area of study is the complexity of state preparation [2], which asks what states can be prepared by quantum computers with bounded resources. The recently proven NLTS theorem [5] states that there exist local Hamiltonians whose low energy states cannot be prepared by constant-depth quantum circuits. More broadly, the longstanding

open question of whether complexity classes QMA and QCMA are equal roughly amounts to asking whether every local Hamiltonian has an efficient classical description of it's ground state [4].[1] Beyond this, the complexity of state preparation has implications in quantum cryptography and physics, with connections to black holes and quantum money [2].

In this work we study unconditional separations in the style of Bravyi, Gosset, and Konig in the inputless setting. Classical *input-independent sampling* problems can be thought of as the classical analog of state-preparation problems, in which the goal is to sample from a fixed n-bit distribution D_n using a classical circuit whose input is fixed to uniformly random bits.[2] While input-dependent problems ask about a classical systems ability to *process* information, input-independent problems instead study what distributions classical systems can *prepare*.

At first glance, it may appear that there is a close connection between input-dependent problems and input-independent sampling problems. If it is hard to map input x to output $f(x)$ in constant-depth, is it also hard to sample from the distribution $(X, f(X))$ where X is uniform? Perhaps surprisingly, the answer to this question is no! To illustrate, consider the parity function, which requires $\Omega(\log n/\log\log(n))$ depth to implement with a classical circuit with unbounded fan-in [15]. Despite this, there is a simple depth-2 classical circuit which maps a random string $r \in \{0,1\}^{n-1}$ to output $(X, \text{parity}(X))$ for uniformly random X. This circuit is easy to describe: simply map input r to output

$$(r_1, r_1 \oplus r_2, r_2 \oplus r_3, \ldots, r_{n-2} \oplus r_{n-1}, r_{n-1})$$

and check that the output distribution has the desired statistics. A similar trick can be used to sample from the distribution $(X, \mathsf{PHP}_n(X))$ where PHP_n is the Parity Halving Problem, a search problem introduced in [26] which separates QNC^0 circuits from constant-depth classical circuits with unbounded fan-in.

Indeed, in contrast to search problems, where lower bounds against constant-depth circuits have a long history [15, 19, 21], lower bounds for input-independent problems have only been developed recently. Particularly relevant to this paper is a breakthrough result of Viola [23] in which he gave the first example of a distribution that could not be sampled by constant-depth classical circuits with bounded fan-in, even up to additive error. (In a follow-up work [24], Viola also gave a distribution that can not be sampled by constant-depth classical circuits with *unbounded* fan-in. While this result is stronger, the techniques used in [24] are less natural in the situation studied here).

A natural question is whether constant-depth quantum circuits can sample from distributions that classical circuits cannot. Indeed, the authors of [9] asked exactly this question:

▶ **Question 1** (From [9]). *Does there exist a family of quantum circuits $\{C_n\}_{n\in\mathbb{N}}$ such that, for each $n \in \mathbb{N}$, any constant-depth classical circuit with bounded fan-in (NC^0) with access to uniformly random bits produces a distribution* far *from the output distribution produced by C_n run on the all-zero state?*

In the question above we understand *close* and *far* in the sense of additive error (or total variation distance). We quickly review the definition of this distance below.

[1] Although proving QMA $\neq$ QCMA would also imply P $\neq$ PSPACE, and so is unlikely without sophisticated new tools.

[2] More formally, the goal, given a family of distributions $\{D_n\}$ that depend only on n, is to produce a family of circuits $\{C_n\}$, each of which samples from the appropriate distribution given random bits as input.

▶ **Definition 2** (Total Variation Distance, Δ). *The Total Variation Distance (or Statistical Distance) between two distributions D_1, D_2 over $\{0,1\}^m$ is*

$$\Delta(D_1, D_2) := \max_{T \subseteq \{0,1\}^m} \left| \Pr[D_1 \in T] - \Pr[D_2 \in T] \right| \tag{1}$$

$$= \frac{1}{2} \sum_{a \in \{0,1\}^m} \left| \Pr[D_1 = a] - \Pr[D_2 = a] \right| \tag{2}$$

In the next section we discuss the main results of this paper, including a positive answer to Question 1 when the number of random inputs given to the classical circuit is bounded.

2 Results

The main result of this paper is the following Theorem.

▶ **Theorem 3.** *For each $\delta \in [0,1)$, there exists a family of distributions $\{D_n\}$ such that for each $n \in \mathbb{N}$, D_n is a distribution over $\{0,1\}^n$ and*

1. *There exists a uniform family of constant-depth quantum circuits $\{C_n\}$ such that for each n, applying C to input $|0^n\rangle$ produces a distribution which has total variation distance at most $1/6 + O(n^{-c})$ from D_n for some $c \in (0,1)$.*
2. *Each classical circuit with fan-in 2 which takes $n + n^\delta$ random bits as input and has total variation distance at most $\frac{1}{2} - \omega(1/\log n)$ from D_n has depth $\Omega(\log \log n)$.*

The distributions D_n constructed are of the form $(X, f(X))$ for a uniformly random bitstring X and function $f : \{0,1\}^{n-1} \to \{0,1\}$.

To provide context for the classical lower bound we note that a uniformly random bitstring has total variation distance $1/2$ from the distribution D_n (or any other distribution of the form $(X, g(X))$ for uniformly random X and function $g : \{0,1\}^{n-1} \to \{0,1\}$) and so the classical lower bound on total variation distance is near-optimal.

Considering the family of constant-depth quantum circuits that approximately produce the distributions $\{D_n\}$, we get the following Corollary, showing the answer to Question 1 is YES provided the number of random bits provided the classical circuit is bounded:

▶ **Corollary 4.** *There exists a uniform family of constant-depth quantum circuits $\{C_n\}$ such that, for each $\delta \in [0,1)$, any classical circuit with fan-in 2 which takes $n + n^\delta$ random bits as input and samples from the n-bit output distribution of C_n to within $1/3 - \omega(1/\log n)$ additive error has depth $\Omega(\log \log n)$.*

While we view Theorem 3 and Corollary 4 as the main results of the paper, we observe that they have some limitations, which we address in part with subsequent theorems. Perhaps most significantly, the quantum circuits C_n constructed in Theorem 3 involve arbitrary constant-sized unitaries. In Appendix A of the full version of the paper we review a standard series of arguments which shows that these unitaries can also be compiled in constant depth by circuits consisting of arbitrary single qubit gates and two-qubit CNOT gates. This shows, in particular, that the quantum circuits $\{C_n\}$ are a *uniform* circuit family.

Additionally, it should be noted that even arbitrary single qubit gates have some capabilities which are beyond the reach of NC^0 circuits with uniformly random input. In particular, applying a single controlled-X rotation to a qubit initially in the $|0\rangle$ state and then measuring in the computational basis results in a random bit sampled from a Bernoulli distribution with arbitrary bias (determined by the extent of the rotation). For most biases, reproducing this bias exactly with an NC^0 given uniformly random input requires super-constant

depth. It seems possible to build on this observation and produce a separation similar to the one appearing in Theorem 3 – indeed, independent from this observation being made here but while we were revising the paper to discuss this issue, this observation was also made formal in [17, Theorem 1.10]. The authors of that paper also show this observation gives a classical-quantum separation that holds even when the number of (uniformly random) input bits provided to the NC^0 circuit is unbounded.

To address this issue, in Appendix C of the full version of this paper we extend the classical lower bound in Theorem 3 to the setting where the classical NC^0 circuit has biased inputs. In this setting we can also prove a lower bound when there is an increased number of biased bits provided as input to the circuit, as long as the overall entropy of the input is not too large – in particular we give a lower bound when the classical circuit has access to $kn + n^\delta$ random bits with each bit drawn from a Bernoulli distribution with entropy $1/k$.

Finally, we reemphasize that the classical lower bound in Theorem 3 only applies to NC^0 circuits which take as input $n + n^\delta$ random bits for some $\delta < 1$. That is, the bound (as well as the generalization in Appendix C of the full version of this paper discussed in the previous paragraph) only applies to NC^0 circuits with access to at most an extra n^δ bits of randomness on top of what is required to sample from the distribution $(X, f(X))$. Because of this restriction we took significant care to construct quantum circuits C_n which only involve n qubits, ensuring a fair comparison. But we also note that subsequent works [24], Viola proves sampling lower bounds which hold against stronger circuit classes, and without the restriction on the number of random input bits. So far we have been unable to adapt those stronger bounds to our setting – this question is discussed in more length in Section 4. As a potential first step in this direction, in Appendix B of the full version of the paper we consider classical circuits with an *unlimited* number of inputs but that have bounded fan-in *and* fan-out. In this setting we show there is a distribution D'_n which can be approximately sampled from by QNC^0 circuits with an advice state, but which cannot be sampled in constant depth classically.

Table 1 Table comparing a few different computational problems with either conditional or unconditional proof of quantum advantage.

Problem	*classical hardness*	*constant depth*	*unconditional*	*input-independent*
Factoring [20]	Poly-time	X[3]	X	X
Sampling depth-3 quantum circuits [22, 1]	Poly-time	✓	X	X
Random Circuit Sampling [7, 3, 6]	Poly-time	✓	X	X
2D-HLF [9]	NC^0	✓	✓	X
This work	NC^0	✓	✓	✓

3 Technical Overview

The distribution used to prove Theorem 3 is a variation of the distribution $(X, \text{majmod}_p(X))$, where the function majmod_p ("Majority mod p") is defined as

$$\text{majmod}_p(x) = \begin{cases} 0 & \text{if } |x| < p/2 \mod p \\ 1 & \text{if } |x| > p/2 \mod p \end{cases} \quad \text{for each } x \in \{0,1\}^{n-1}, \text{ and prime } p. \tag{3}$$

[3] Factoring *can* accomplished in logarithmic depth on a quantum computer with polynomial time classical post-processing [11], or in constant-depth on quantum computer with unbounded fanout gates [16] or intermediate measurements [10], again with classical post-processing.

Viola introduced the majmod_p function in [23]. In the same paper he showed a hardness of sampling result for the distribution $(X, \text{majmod}_p(X))$ similar to Item 2 of Theorem 3.

To illustrate some key ideas used in the proof of Theorem 3, we first sketch the proof of a weaker sampling separation which holds when QNC^0 circuits take as input a GHZ state $|\text{GHZ}_n\rangle = 1/\sqrt{2}\,(|0^n\rangle + |1^n\rangle)$. We refer to these circuits as QNC^0 states with GHZ *advice*. The statement of this separation is as follows.

▶ **Theorem 5** (Separation with GHZ advice). *For each $n \in \mathbb{N}$, and $\delta \in [0, 1)$, there exists a prime p such that*

1. *There exists a constant-depth quantum circuit that takes the GHZ_n state as input and produces a distribution which has total variation distance at most $1/6 + O(n^{-c})$ from $(X, majmod_p(X) \oplus parity(X))$ for some $c \in (0, 1)$.*
2. *Each classical circuit with bounded fan-in which takes $n + n^\delta$ random bits as input and has total variation distance at most $\frac{1}{2} - \omega(1/\log n)$ from $(X, majmod_p(X) \oplus parity(X))$ has depth at least $\Omega(\log\log(n))$.*

We sketch the proof of Items (1) and (2) of the above theorem in the next two subsections.

3.1 Approximately sampling from $(X, \text{majmod}_p(X) \oplus \text{parity}(X))$ with GHZ advice

We first introduce some non-unitary "self-controlled" rotation operations, then describe a circuit that approximately samples from $(X, \text{majmod}_p(X) \oplus \text{parity}(X))$ using those operations. We close this subsection by showing that it is possible to approximate those non-unitary operations with unitary gates.

3.1.1 Introducing non-unitary gates

We introduce the following single-qubit non-unitary operator.

$$A_\theta := |0\rangle\langle 0| + e^{-i\theta X}\,|1\rangle\langle 1|\,, \qquad \theta \in \mathbb{R} \tag{4}$$

It is straightforward to see that A_θ is linear, but not unitary. A_θ can be interpreted as a "self-controlled" X rotation gate. That is, applied to the $|0\rangle$ state, it acts as the identity, and on the $|1\rangle$ state, an $e^{i\theta X}$ is applied. For this reason, it is convenient to draw the gate and its adjoint as

$A_\theta =$ 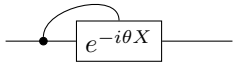 $A_\theta^\dagger =$ 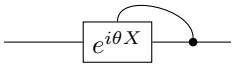

Upon post-selection on the output of $A_\theta^\dagger$ in the computational basis, we can simplify these operations as follows:

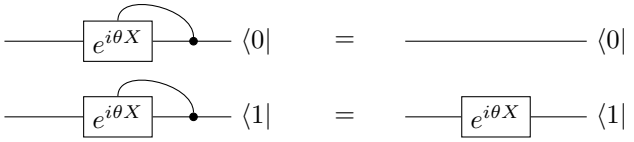

(5)

3.1.2 Approximate sampling with a non-unitary circuit

Next, we illustrate how these non-unitary gates and GHZ advice can be used together to produce an n-bit distribution where the first $(n-1)$ bits are uniformly random, and the final bit is a function of the Hamming weight of the first $(n-1)$ which approximates $(X, \text{majmod}_p(X) \oplus \text{parity}(X))$. We proceed by walking through an example on $n = 4$ qubits.

Starting with the $|\text{GHZ}_4\rangle$ state, we apply a Hadamard gate to each qubit, and then apply our "self-controlled" rotation gates $A_\theta^\dagger$ to all but the last qubit as shown on the left hand side below. To analyze the output distribution of this circuit we proceed with a series of circuit identities. The first is the fact that $H^{\otimes n}$ maps the state $|\text{GHZ}_n\rangle$ to the $|\text{EVEN}_n\rangle$ state. Here, the EVEN_n state denotes the uniform superposition over all even n-bit strings.

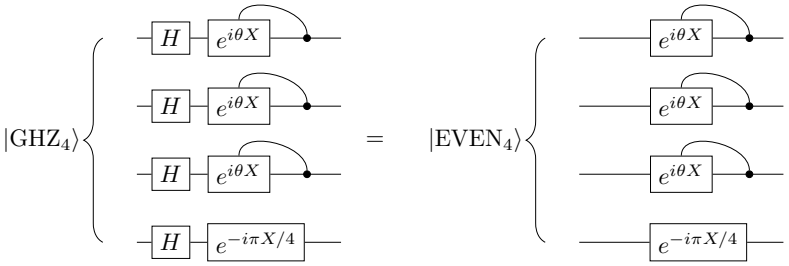

Suppose we measure the first 3 qubits in the $\{|0\rangle, |1\rangle\}$ basis, getting outcomes $x_1, x_2, x_3 \in \{0, 1\}$. Using the circuit identities in Equation (5), we have

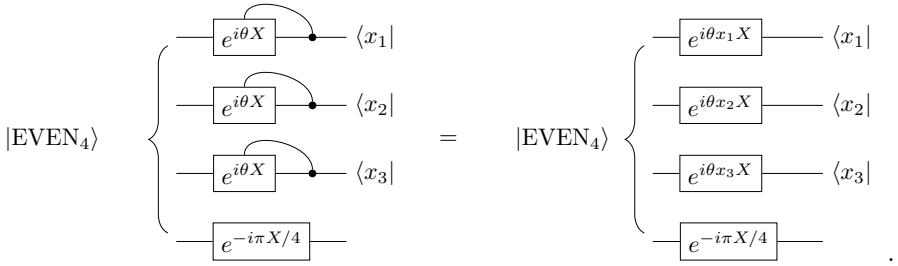

Next, we observe that if you apply the Pauli-X operator to any single qubit of the $|\text{EVEN}_n\rangle$ state, it becomes the $|\text{ODD}_n\rangle$ state. Therefore, it has the same effect as if we instead applied X to, say, the last qubit. The same is true for X-rotation gates, $\exp(i\theta X)$, so we can push all of our gates down to the last qubit.

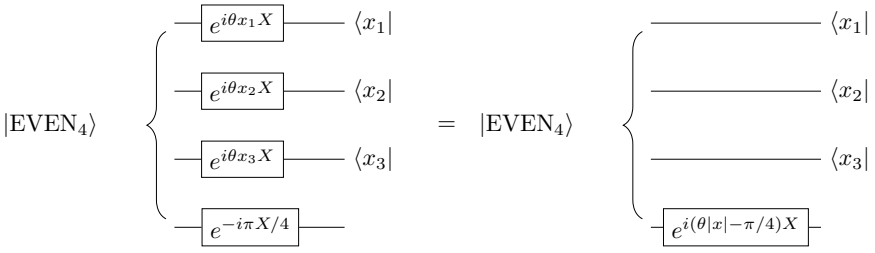

Finally, we note that the $|\text{EVEN}_n\rangle$ state can be constructed by first initializing the first $n-1$ qubits to the $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ state, and the final qubit in the $|0\rangle$ state, and subsequently computing the parity of the first $n-1$ qubits into the final register[4].

[4] Parity can be implemented with a sequence of CNOT gates: $\prod_{i=1}^{n-1} \text{CNOT}_{i,n} |x\rangle |0\rangle = |x\rangle |\text{parity}(x)\rangle$ for each $x \in \{0,1\}^{n-1}$

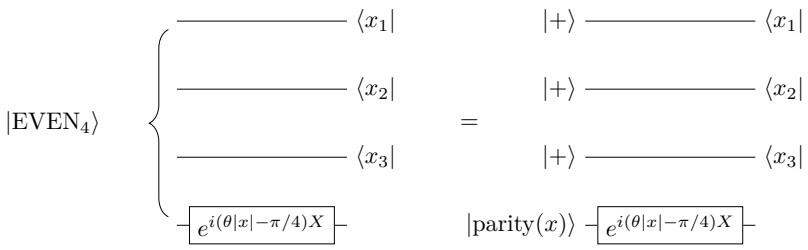

And so our measurement outcomes on the first $n-1$ bits are uniformly random, as desired. As for the last qubit – let $b \in \{0,1\}$ be the outcome after measuring the last qubit in the standard basis. Then we have that

$$\Pr[b = \text{parity}(x)] = \cos^2(\theta|x| - \pi/4)\,. \tag{6}$$

This function is periodic in the Hamming weight of x, with periodicity π/θ. If we set $\theta = \pi/p$ then the output bit b approximately correlates with $\text{majmod}_p(x) \oplus \text{parity}(x)$, as can be verified analytically or by inspection of the following figures:

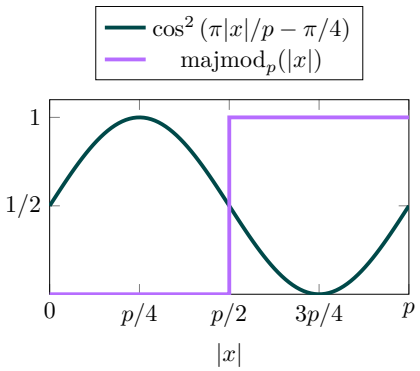

(a) Inverse correlation of $\cos^2(\pi|x|/p - \pi/4)$ and $\text{majmod}_p(x)$.

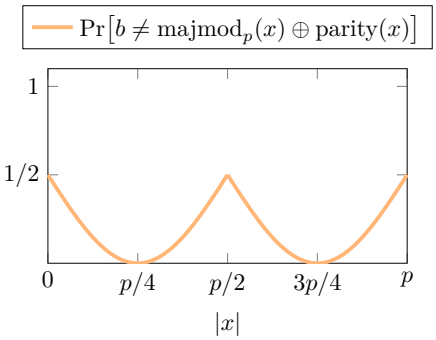

(b) Probability that $b \neq \text{majmod}_p(x) + \text{parity}(x)$ provided $\theta = \pi/p$.

Figure 1 Plots displaying the correlation of the bit b in the circuits above and the function $\text{majmod}_p(x) + \text{parity}(x)$ provided $\theta = \pi/p$.

A straightforward calculation (Lemma 8 in the full version of the paper) then shows the output distribution of the non-unitary circuit described above is close in total variation distance to the distribution $(X, \text{majmod}_p(X) \oplus \text{parity}(X))$.

3.1.3 Converting to Unitary

Our non-unitary circuits are helpful for initial circuit design, but we need to somehow port them back over to be unitary – while maintaining their low-depth. To this end, we make use of the following two insights to construct a unitary circuit such that the output distribution is very close to that of the non-unitary circuit.

1. We do not need to find a unitary that is close to the circuit (in fact, this is likely not possible). It is sufficient to instead find a unitary that has the same behavior with respect to its action on the GHZ state.
2. We introduce a multi-qubit non-unitary gate $A_{m,\theta}$ acting on m-qubits that has the same action as $A_\theta^{\otimes m}$ when applied to the GHZ state, and becomes closer to unitary as m increases.

In Section 5 of the full version of the paper we make the outline above rigorous to construct a quantum circuit that, with advice, samples approximately from the distribution $(X, \text{majmod}_p(X) \oplus \text{parity}(X))$, where $X \sim \text{Unif}\left(\{0,1\}^{n-1}\right)$.

3.2 Classical circuit lower bound for $(X, \text{majmod}_p(X) + \text{parity}(X))$

The proof of a classical lower bound for the distribution $(X, \text{majmod}_p(X) + \text{parity}(X))$ closely follows Viola's techniques in [23]. Rather than explicitly lower bounding classical circuit depth, Viola proves lower bounds for the *locality* of functions. To illustrate the relationship between locality and circuit depth let $f : \{0,1\}^\ell \to \{0,1\}^n$ be a function implemented by a classical circuit with bounded fan-in. We say that f is ℓ_f-local if, for each $i \in [n]$, the i-th output bit of $f(u)$ depends on at most ℓ_f bits of the input u. Any circuit with depth d and fan-in χ will have locality at most χ^d. And so, to prove a circuit lower bound of $\Omega(\log \log n)$ for sampling from the distribution $(X, \text{majmod}_p \oplus \text{parity}(X))$ it suffices to prove that there exists some $k > 0$ such that any function with locality at most $\Omega(\log^k n)$ cannot sample from the distribution $(X, \text{majmod}_p \oplus \text{parity}(X))$ given access to uniformly random bits as input.

Both our proof of sampling hardness for $(X, \text{majmod}_p(X) \oplus \text{parity}(X))$ and Viola's original proof of hardness for $(X, \text{majmod}_p(X))$ begin with the observation that for any ℓ_f-local function $f : \{0,1\}^\ell \to \{0,1\}^n$ there exists a partition of the input $u = (x, y)$ and a permutation of output bits of $f(x, y)$ such that:[5]

$$f(x, y) = g_1(x_1, y) \circ g_2(x_2, y) \circ \cdots \circ g_s(x_s, y) \circ h(y), \tag{7}$$

where each $g_i(x_i, y)$ is a subset (or "block") of the output bits that are completely determined by y and a single bit of x, and $s = \Omega(n/\ell_f^2)$. Therefore, if we fix the input bits y, each of the blocks g_i are independent. Let $z \in \{0,1\}^{n-1}$ be the first $n-1$ outputs of $f(x, y)$ and let b be the final output bit. We can also assume without loss of generality (by absorbing at most one g_i into h) that the last output bit is not permuted, so b only depends on y. In order for the function f to sample from the correct distribution the output bits z must be uniformly distributed and, for every input (x, y), we must have $\text{majmod}_p(z) \oplus \text{parity}(z) = b$. After fixing the input bits y, the Hamming weight of z is a sum of independent random variables but b is fixed. Then (still following Viola) we show that if many of these independent variables are fixed the output distribution of z will not have sufficiently high entropy. Alternatively, if they are unfixed, the condition $\text{majmod}_p(z) \oplus \text{parity}(z) = b$ is unlikely to be satisfied. Making these observations formal completes the proof. The full details of this argument are given in Section 7 of the full version of the paper.

3.3 Removing the GHZ advice

To complete the proof sketch of Theorem 3 we need to extend a proof of Theorem 5 to give a sampling separation *without* GHZ advice. To do this, we replace the GHZ state in the quantum circuit used to prove Theorem 5 with a "Poor-Man's GHZ state" (introduced in [26]) defined over a binary tree $\mathcal{B}$. An n qubit Poor-Man's GHZ state can be prepared by a constant-depth circuit acting on $2n-1$ qubits followed by a measurement of $n-1$ auxiliary qubits. The remaining state is equivalent to the GHZ state with some Pauli terms applied to it. We can determine which Pauli operations will "correct" the state back to the GHZ state

[5] We use "$\circ$" to denote concatenation.

as a function of our measurement outcomes. However, determining these corrections requires $\Omega(\log n)$ depth – which we cannot afford in this shallow circuit setting. Instead, we absorb the Pauli corrections into the definition of the target distribution.

The result is a circuit that samples approximately from a modified version of the $(X, \text{majmod}_p(X) \oplus \text{parity}(X))$ distribution. Specifically, the new circuit (including the measured auxiliary qubits) approximately samples from a distribution of the form $(X, \text{MM}_p(S_X) \oplus \text{parity}(X))$ where

$$\text{MM}_p(j) := \begin{cases} 0 & \text{if } j < p/2 \mod p \\ 1 & \text{if } j > p/2 \mod p \end{cases} \qquad \text{for } j \in \mathbb{Z}. \tag{8}$$

and S_X is a sum of terms that depends on output bits $X \in \{0,1\}^{n-1}$. In particular, S_X is a weighted sum of the bits of X where the sign of the weight for each bit X_i may depend on many other output variables. In the body of the paper we introduce the notation $\text{MM}_p(S_X) \oplus \text{parity}(X) := \text{pmmajmod}_p(X)$ to describe this new function.

Unlike the function $\text{majmod}_p(X) \oplus \text{parity}(X)$, the function $\text{pmmajmod}_p(X)$ does not just depend on the Hamming weight of X. This introduces a complication when trying to show the classical hardness of sampling using Viola's previously discussed lower bounding technique, since this technique relied on the fact that the Hamming weight of X could be written as a sum of the Hamming weights of disjoint blocks g_i of output bits (and that these blocks became independent after fixing enough input bits). To get around this we show that, after fixing additional input bits (and hence some output bits), we can find blocks of unfixed output bits which each depend on disjoint single input bits and which contribute to disjoint terms in the sum S_X. After showing this additional detail, the proof of the lower bound proceeds similarly to Viola's. Although perhaps conceptually straightforward, this argument is mathematically delicate, and relies on careful counting related to the binary tree layout used to construct the poor man's GHZ state. Details of the new circuit construction (with the "Poor Man's GHZ state") are given in Section 6 of the full version of the paper, while details of the classical lower bound for this new function are given in Section 8.

4 Discussion and Open Problems

Our results show that QNC^0 circuits can sample from distributions that NC^0 circuits cannot. Below we list a few ways in which we think these results could potentially be extended.

- In an experiment with the goal of demonstrating quantum advantage, one would like to not just construct a QNC^0 circuit that samples from a distribution which NC^0 circuits cannot, but also *verify* that the distribution sampled from is indeed hard to sample from classically. How many samples are needed for this verification? Can the circuit be modified to make the verification easier? We point out here that the constant total variation distance in Corollary 4 means that only a few samples are needed to verify that the distribution produced by the described quantum circuit is not produced by a fixed NC^0 circuit, for any specific choice of circuit. However, ruling out *all* distributions producible by NC^0 circuits is a harder task.
- The procedure described in Appendix A of the full version of the paper for compiling the $U_{m,\theta}$ unitary is unlikely to produce an "optimal" compilation. With careful thought it may be possible to find a more natural compilation technique that produces $U_{m,\theta}$ gates while requiring many fewer elementary gates. Finding such a compilation would likely make an experimental implementation of the circuits described in this paper much more feasible.

- Can we get rid of the limitation on the number of inputs to the classical circuit? In Appendix B of the full version of the paper we make some progress in this direction. We consider classical circuits with an *unlimited* number of inputs but that have bounded fan-in *and* fan-out. We show that such classical circuits of depth $o(\log \log n)$ produce distributions far from $D = (X, \text{majmod}_p(X))$. Whereas, we show in Section 5 of the full version of the paper that constant-depth quantum circuits with bounded fan-in and fan-out, when given a GHZ advice state, can sample close to D.
- Can we prove an input-independent sampling separation between QNC^0 and AC^0 circuits? Notably, in [24], Viola proves certain distributions cannot be produced by AC^0 circuits. Can these techniques be extended to QNC^0 circuits? If so, we would have a novel technique for lower-bounding the circuit complexity of quantum states. If not, we should be able to find a QNC^0 circuit that samples from one of these distributions, producing the desired sampling separation.

References

1 Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461(2063):3473–3482, 2005.

2 Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016. arXiv:1607.05256.

3 Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. *arXiv preprint arXiv:1612.05903*, 2016. arXiv:1612.05903.

4 Dorit Aharonov and Tomer Naveh. Quantum np-a survey. *arXiv preprint quant-ph/0210077*, 2002.

5 Anurag Anshu, Nikolas Breuckmann, and Chinmay Nirkhe. Nlts hamiltonians from good quantum codes. *arXiv preprint arXiv:2206.13228*, 2022. doi:10.48550/arXiv.2206.13228.

6 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.

7 Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. Quantum supremacy and the complexity of random circuit sampling. *arXiv preprint arXiv:1803.04402*, 2018. arXiv:1803.04402.

8 Sergey Bravyi, David Gosset, Robert Koenig, and Marco Tomamichel. Quantum advantage with noisy shallow circuits. *Nature Physics*, 16(10):1040–1045, 2020.

9 Sergey Bravyi, David Gosset, and Robert König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018.

10 Dan Browne, Elham Kashefi, and Simon Perdrix. Computational depth complexity of measurement-based quantum computation. In *Conference on Quantum Computation, Communication, and Cryptography*, pages 35–46. Springer, 2010. doi:10.1007/978-3-642-18073-6_4.

11 Richard Cleve and John Watrous. Fast parallel circuits for the quantum fourier transform. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 526–536. IEEE, 2000. doi:10.1109/SFCS.2000.892140.

12 Iulia Georgescu. How the bell tests changed quantum physics. *Nature Reviews Physics*, 3(10):674–676, 2021.

13 Daniel Grier and Luke Schaeffer. Interactive shallow clifford circuits: Quantum advantage against nc^1 and beyond. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 875–888, 2020. doi:10.1145/3357713.3384332.

14 Aram W Harrow and Ashley Montanaro. Quantum computational supremacy. *Nature*, 549(7671):203–209, 2017. doi:10.1038/NATURE23458.

15 Johan Torkel Håstad. *Computational limitations for small-depth circuits*. MIT press, 1987.

16 Peter Høyer and Robert Špalek. Quantum fan-out is powerful. *Theory of computing*, 1(1):81–103, 2005. doi:10.4086/TOC.2005.V001A005.

17 Daniel M Kane, Anthony Ostuni, and Kewen Wu. Locality bounds for sampling hamming slices. *arXiv preprint arXiv:2402.14278*, 2024. doi:10.48550/arXiv.2402.14278.

18 John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018. doi:10.22331/Q-2018-08-06-79.

19 Alexander A Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.

20 Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999. doi:10.1137/S0036144598347011.

21 Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 77–82, 1987. doi:10.1145/28395.28404.

22 Barbara M. Terhal and David P. DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and arthur-merlin games, 2002. arXiv:quant-ph/0205133.

23 Emanuele Viola. The complexity of distributions. *SIAM Journal on Computing*, 41(1):191–218, 2012. doi:10.1137/100814998.

24 Emanuele Viola. Extractors for circuit sources. *SIAM Journal on Computing*, 43(2):655–672, 2014. doi:10.1137/11085983X.

25 John Watrous. Quantum computational complexity. *arXiv preprint arXiv:0804.3401*, 2008.

26 Adam Bene Watts, Robin Kothari, Luke Schaeffer, and Avishay Tal. Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 515–526, 2019. doi:10.1145/3313276.3316404.