On the Pauli Spectrum of QAC0

Shivam Nadimpalli
Columbia University
New York, USA
sn2855@columbia.edu

Natalie Parham
Columbia University
New York, USA
natalie@cs.columbia.edu

Francisca Vasconcelos
University of California
Berkeley, USA
francisca@berkeley.edu

Henry Yuen
Columbia University
New York, USA
hyuen@cs.columbia.edu

ABSTRACT

The circuit class QAC^0 was introduced by Moore (1999) as a model for constant depth quantum circuits where the gate set includes many-qubit Toffoli gates. Proving lower bounds against such circuits is a longstanding challenge in quantum circuit complexity; in particular, showing that polynomial-size QAC^0 cannot compute the parity function has remained an open question for over 20 years.

In this work, we identify a notion of the *Pauli spectrum* of QAC^0 circuits, which can be viewed as the quantum analogue of the Fourier spectrum of classical AC^0 circuits. We conjecture that the Pauli spectrum of QAC^0 circuits satisfies *low-degree concentration*, in analogy to the famous Linial, Mansour, Nisan (LMN) theorem on the low-degree Fourier concentration of AC^0 circuits. If true, this conjecture immediately implies that polynomial-size QAC^0 circuits cannot compute parity.

We prove this conjecture for the class of depth-d, polynomial-size QAC^0 circuits with at most $n^{O(1/d)}$ auxiliary qubits. We obtain new circuit lower bounds and learning results as applications: this class of circuits cannot correctly compute

- the n-bit parity function on more than $(\frac{1}{2} + 2^{-\Omega(n^{1/d})})$-fraction of inputs, and
- the n-bit majority function on more than $(\frac{1}{2} + O(n^{-1/4}))$-fraction of inputs.

Additionally we show that this class of QAC^0 circuits with limited auxiliary qubits can be learned with quasipolynomial sample complexity, giving the first learning result for QAC^0 circuits.

More broadly, our results add evidence that "Pauli-analytic" techniques can be a powerful tool in studying quantum circuits.

CCS CONCEPTS

• **Theory of computation** → **Quantum complexity theory**; **Circuit complexity**.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649662

KEYWORDS

QAC0, analysis of Boolean functions, quantum circuit complexity

ACM Reference Format:
Shivam Nadimpalli, Natalie Parham, Francisca Vasconcelos, and Henry Yuen. 2024. On the Pauli Spectrum of QAC0. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3618260.3649662

1 INTRODUCTION

The Fourier spectrum of a Boolean function provides highly useful information about its complexity. For example, the celebrated result of Linial, Mansour, and Nisan [33] shows that polynomial-size AC^0 circuits give rise to functions whose Fourier spectra obey *low-degree concentration*; in other words, they are approximately low-degree polynomials. Since then, Fourier analytic techniques have played an essential role in breakthroughs in learning theory [15, 29, 32], pseudorandomness [8, 10, 11], property testing [5, 17, 31], classical-quantum separations [40], and more.

Are there analogous analytic techniques for studying models of quantum computation? A natural quantum generalization of the Fourier spectrum of a Boolean function is the *Pauli spectrum* of a many-qubit operator. Recall that the set of n-qubit Pauli operators, $\mathcal{P}_n := \{I, X, Y, Z\}^{\otimes n}$, forms an orthonormal basis for the space of $2^n \times 2^n$ complex matrices, with respect to the (normalized) Hilbert–Schmidt inner product. Consequently, any n-qubit operator A can be decomposed as $A = \sum_{P \in \mathcal{P}_n} \widehat{A}(P)\, P$, analogous to the Fourier expansion of a Boolean function. Our paper is motivated by the following question:

When does the Pauli spectrum reveal useful information about the complexity of a quantum computation?

Analyzing the Pauli spectrum of quantum operations has been a fruitful endeavor in recent years, leading to structural results about so-called quantum Boolean functions [35, 43], learning algorithms for quantum dynamics [26, 46, 50], property testing of quantum operations [3, 13, 51], and classical simulations of noisy quantum circuits [1]. Although each result uses a slightly different notion of Pauli decomposition, most of them focus on the Pauli spectrum of unitary operators. However, it is unclear how this notion connects with the complexity of the unitary operator.

An Illustrative Example. One might have hoped for the following Linial-Mansour-Nisan-style low-degree concentration statement: if a unitary U is computable by some simple circuit (for some notion

of "simple"), then most of its "Pauli mass" should concentrate on its low-degree part. The *degree* of a Pauli tensor $P \in \mathcal{P}_n$, denoted $|P|$, is the number of qubits on which it acts non-trivially (i.e. the number of non-identity components). Unfortunately, however, such notions of low-degree Pauli concentration break down for even the simplest unitaries: consider the unitary $U = X^{\otimes n}$, which can be implemented by a single layer of single-qubit gates (see Figure 1). Clearly the Pauli mass of U is concentrated on a single degree-n coefficient, yet this unitary U is computed by an extremely simple circuit.

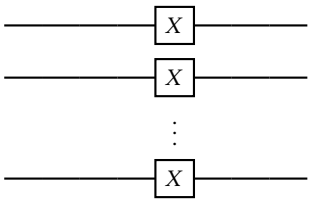

Figure 1: A simple unitary that does not satisfy low-degree Pauli concentration in the naive sense.

Note, however, that in this example there is an incongruity between the classical and quantum settings. A classical Boolean function $f : \{0,1\}^n \to \{0,1\}$ only has one output bit, whereas the output of unitary U is n qubits. This leads us to restrict our focus to quantum circuits where there is a designated "target" qubit in the output. In the aforementioned example, while unitary U has degree-n, the target qubit is not influenced by any other qubits, meaning U should "morally" have degree-1. This suggests that instead of analyzing the unitary operator corresponding to the circuit, one should analyze the *quantum channel* that comes from applying the circuit and then tracing out all qubits except for the target qubit.

A New Notion of Pauli Spectrum. We introduce a different notion of Pauli decomposition: The Pauli spectrum of a quantum channel $\mathcal{E}$ mapping n qubits to ℓ qubits is the set of Pauli coefficients of its $(n+\ell)$-qubit *Choi representation* $\Phi_{\mathcal{E}} := (I^{\otimes n} \otimes \mathcal{E})(|\mathrm{EPR}_n\rangle\langle\mathrm{EPR}_n|)$ which is the channel applied to half of un-normalized n-qubit Bell state $|\mathrm{EPR}_n\rangle := \sum_{x\in\{0,1\}^n} |x\rangle \otimes |x\rangle$ (see the full version of the paper for formal definitions of Choi matrices and their Pauli spectrum).

We show that this definition of Pauli spectrum connects with computational complexity much more closely. In particular we illustrate the definition's usefulness by studying the Pauli spectrum of QAC^0 circuits. These are a model of shallow quantum circuits that, in addition to single-qubit gates, can include wide Toffoli gates acting on arbitrarily many qubits [36]. This is a quantum analogue of the classical circuit model AC^0, and represents the frontier of quantum circuit complexity: although we know that polynomial-size AC^0 circuits cannot compute parity [18, 22], proving the same for QAC^0 circuits remains a longstanding open problem [4, 16, 25, 36, 39, 42].

Our Results, In a Nutshell. Our main technical result, at a high level, is that polynomial-size, single-qubit-output QAC^0 circuits that use few auxiliary qubits must have Pauli spectrum that is highly concentrated on low-degree coefficients. This is a new structural result about QAC^0 circuits (with limited auxiliary qubits) that immediately yields average-case circuit lower bounds: we show that such circuits cannot compute the parity or majority functions, even approximately (see Section 3 for detailed theorem statements).

This raises the intriguing question of whether low-degree concentration holds for *all* polynomial-size QAC^0 circuits, not just ones with few auxiliary qubits. We conjecture that this is indeed true (see Conjecture 1 for a formal statement). This would be directly analogous to the Linial-Mansour-Nisan theorem about the low-degree concentration of AC^0 circuits [33], and would immediately show that parity is not computable by polynomial-size QAC^0 circuits, resolving the central open question posed in Moore's 1999 paper introducing the QAC circuit model in the first place [36].

Finally, we show that low-degree concentration of the Pauli spectrum of quantum channels yields sample-efficient *learning algorithms* for those channels (see Theorem 3 for a more detailed theorem statement). This also directly corresponds to the learning result of [33] who show that low-degree concentrated Boolean functions can be learned with quasipolynomial complexity. Our results directly imply that QAC^0 circuits with few auxiliary qubits can be learned using quasipolynomial sample complexity, and if the conjectured low-degree concentration of QAC^0 holds, then *all* polynomial-size QAC^0 circuits are sample-efficiently learnable.

Although we weren't able to prove "quantum LMN," we believe that the conjecture provides tantalizing motivation for studying the analytic structure of QAC^0 circuits, and for further investigating this notion of Pauli spectrum more broadly. The analogy with classical Fourier analysis of Boolean functions appears quite strong; it will be interesting to discover how far the analogy goes.

Before explaining our results in more detail, we give a brief overview of QAC^0 circuits.

2 QAC^0 CIRCUITS

The QAC^0 circuit model consists of constant-depth quantum circuits with arbitrary single-qubit gates and arbitrary-width Toffoli gates, which implement the unitary transformation

$$|x_1, \ldots, x_n, b\rangle \mapsto |x_1, \ldots, x_n, b \oplus \bigwedge_i x_i\rangle .$$

QAC^0 and related models were first introduced by Moore [36] to explore natural quantum analogues of classical circuit classes such as NC^0, AC^0, and $\mathrm{AC}^0[q]$, which are well-studied models of shallow circuits in classical complexity theory.

Aside from being a natural generalization of a classical circuit model, QAC^0 also gives a clean theoretical model with which to study the power of many-qubit operations in quantum computations. Recently there has been increasing motivation for understanding the power of short-depth quantum computations with many-qubit or non-local operations. Some experimental platforms are beginning to realize controllable operations that can affect many qubits at once; examples include analog simulators [6], ion traps [20, 21], and superconducting qubit platforms that have mid-circuit measurements [44].

Parity versus QAC^0. Already in his 1999 paper [36], Moore posed the question of whether the n-bit parity function can be computed in QAC^0. Recall that computing parity is equivalent to computing

Table 1: Comparison with the work of Linial, Mansour, and Nisan [33]; see Section 4 for a detailed discussion. In both the classical and the quantum settings, we denote the number of inputs as n, number of gates as s, and the depth of the circuit as d. In the QAC^0 learning setting, we assume that $a \leq O(\mathrm{polylog}(n))$.

	AC^0 Circuits (Boolean function f)	QAC^0 Circuits (Choi representation Φ)
Fourier Basis	$\chi_S(x) \in \left\{\prod_{i\in S} x_i\right\}_{S\subseteq[n]}$ (Parity functions)	$P \in \{I,X,Y,Z\}^{\otimes n}$ (Pauli tensors)
Decomposition	$f(x) = \sum_{S\subseteq[n]} \widehat{f}(S)\chi_S(x)$ (Fourier decomposition)	$\Phi = \sum_{P\in\mathcal{P}_n} \widehat{\Phi}(P)\cdot P$ (Pauli decomposition)
Spectral Concentration	$\mathbf{W}^{>k}[f] \leq s\cdot 2^{-\Omega\left(k^{1/d}\right)}$ (Lemma 7 of [33])	$\mathbf{W}^{>k}[\Phi] \leq s^2\cdot 2^{-\Omega(k^{1/d}-a)}$ (Theorem 1)
Correlation with Parity	$\frac{1}{2} + s\cdot 2^{-\Theta\left(n^{1/d}\right)}$ (Implicit in [33])	$\frac{1}{2} + s\cdot 2^{-\Omega(n^{1/d}-a)}$ (Theorem 2)
Learnability	Quasipolynomial sample (Section 4 of [33])	quasipolynomial samples (Theorem 3)

fanout, i.e. the unitary operation

$$|b, x_1, \ldots, x_n\rangle \mapsto |b, x_1 \oplus b, \ldots, x_n \oplus b\rangle\,,$$

up to conjugation by Hadamard gates [36]. Consequently if QAC^0 circuits could compute parity, then this would imply that QAC^0 would be remarkably powerful: among other things, they would be capable of generating GHZ states, computing the MOD_q function for all q, computing the parity function, and performing phase estimation and approximate quantum Fourier transforms [4, 36].

Despite being open for more than twenty years, this question has seen limited progress. Fang et al. [16] showed that QAC circuits with sublinear auxiliary qubits cannot compute parity, and more recently, Padé et al. [39] showed that depth-2 QAC circuits with arbitrary auxiliary qubits cannot compute parity. Rosenthal [42] proved that parity can be approximated exponentially well by QAC^0 circuits that have exponentially many auxiliary qubits (it was not known before whether QAC^0 circuits of *any* size could approximate parity). Rosenthal also proved that certain classes of QAC^0 circuits require exponential size to approximate parity; however, extending these lower bounds to general QAC^0 circuits seems challenging. (See Section 5 for more details about these prior works.)

Although lower bounds against classical AC^0 are considered one of the great successes of complexity theory [22, 41, 47, 52], it is far from clear how to extend those techniques (such as switching lemmas) to the setting of QAC^0. Furthermore, it is unclear whether QAC^0 lower bounds imply AC^0 lower bounds: we do not know if QAC^0 circuits can even simulate classical AC^0 circuits. Even though QAC^0 circuits appear quite weak (especially for classical computation), lower bounds have been difficult to obtain.

Going Beyond Lightcones. On the other hand, if we restrict ourselves to constant-depth quantum circuits with only two-qubit gates (known as QNC^0 circuits), it is comparatively much easier to prove limitations. For example, such circuits cannot prepare states with long-range entanglement (like the many-qubit GHZ state or states with topological order) [2, 14, 49]. At the heart of all QNC^0 lower bound techniques, is the "lightcone argument"— the observation that each output qubit can only be affected by a small number of input qubits since there are only a few layers of small gates (see Figure 2). This argument immediately breaks when either the circuit has logarithmic depth, or large many-qubit gates. Thus, any effort to prove lower bounds against more general circuits calls for novel techniques *beyond lightcones*. QAC^0 circuits are at the frontier of this boundary and thus an attractive point of attack for developing new circuit lower bound techniques.

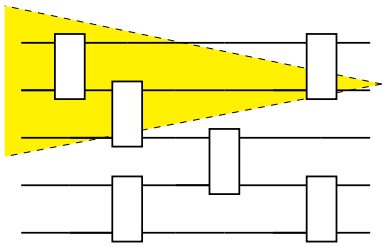

Figure 2: A (backwards) lightcone of a qubit.

3 OUR RESULTS

We show that the spectral properties – with our notion of Pauli spectrum – of QAC^0 circuits with limited auxiliary qubits resemble those of classical AC^0 circuits; this in turn allows us to derive circuit lower bound and learning conclusions that are analogous to those of [33] (see Table 1 for a comparison). In particular, our main technical result is the following bound on the Pauli coefficients of a QAC^0 circuit. For definitions of Choi representations, Pauli coefficients, etc, we refer the reader to full version of the paper. We also defer formal statements of our results to the full version of the paper and instead give informal statements below for ease of exposition:

Theorem 1. Suppose a n to 1-qubit channel $\mathcal{E}$ is computed by a depth-d QAC^0 circuit with a auxiliary qubits. Writing $\Phi_{\mathcal{E}}$ for the *Choi representation* of $\mathcal{E}$, we have

$$\sum_{P\in\mathcal{P}_{n+1}:|P|>k} \left|\widehat{\Phi_{\mathcal{E}}}(P)\right|^2 \leq 2^{-\Omega\left(k^{1/d}-a\right)}$$

where $\{\widehat{\Phi_{\mathcal{E}}}(P)\}$ is the collection of *Pauli coefficients* of $\Phi_{\mathcal{E}}$.

Theorem 1 can be interpreted as saying that most of the "mass" of the Pauli decomposition of the Choi representation of $\mathcal{E}$ lies on Pauli tensors with few non-identity components. As a consequence of Theorem 1, we obtain circuit lower bounds and learning algorithms for QAC^0 circuits.

Theorem 2. Suppose Q is a QAC circuit with depth $d = O(1)$ and at most $\frac{1}{2}\cdot n^{1/d}$ auxiliary qubits. Let $Q(x) \in \{0,1\}$ denote the random measurement outcome in the computational basis of a single output qubit of Q on input $|x\rangle$.

- Q cannot approximate the n-bit parity function $\mathrm{Parity}_n(x) = \sum_{i=1}^n x_i \mod 2$, i.e.

$$\Pr_{x\sim\{0,1\}^n}[Q(x) = \mathrm{Parity}_n(x)] \le \frac{1}{2} + 2^{-\Omega\left(n^{1/d}\right)}.$$

- Q cannot approximate the n-bit majority function $\mathrm{Maj}_n(x) = \mathbf{1}\left\{\sum_{i=1}^n x_i \ge n/2\right\}$, i.e.

$$\Pr_{x\sim\{0,1\}^n}\left[Q(x) = \mathrm{Majority}_n(x)\right] \le \frac{1}{2} + O\left(n^{-1/4}\right).$$

We point out that our QAC^0 lower bounds in Theorem 2 are *average-case*: the circuits fail to approximate parity or majority. The only previously known average-case lower bounds for parity were shown by Rosenthal [42]. Notably, he showed an average-case bound against depth-2 QAC circuits and a size lower bound of $\Omega(n/d)$ for depth d circuits. For a more detailed comparison between our lower bounds and previously established QAC^0 lower bounds, see Section 5.

As a consequence of Theorem 1, we also obtain an algorithm for learning QAC^0 circuits:

Theorem 3. Let $\mathcal{E}$ be a channel computed by a depth-d QAC^0 circuit on n input qubits with $\mathrm{polylog}(n)$ auxiliary qubits. Then for $\delta > 0$ and $\epsilon > \frac{1}{\mathrm{poly}(n)}$, it is possible to learn a channel $\widetilde{\mathcal{E}}$ satisfying

$$\sum_{P\in\mathcal{P}_{n+1}} \left|\widehat{\Phi_{\mathcal{E}}}(P) - \widehat{\Phi_{\widetilde{\mathcal{E}}}}(P)\right|^2 \le \epsilon, \tag{1}$$

with probability $1-\delta$ using $n^{\mathrm{polylog}(n)}\log\left(1/\delta\right)$ copies of the Choi state $\frac{1}{N}\Phi_{\mathcal{E}}$. Here, $\Phi_{\mathcal{E}}$ and $\Phi_{\widetilde{\mathcal{E}}}$ are the Choi representations of $\mathcal{E}$ and $\widetilde{\mathcal{E}}$ respectively. In the special case where $\mathcal{E}$ computes a Boolean function $f : \{0,1\}^n \to \{0,1\}$, the learned channel $\widetilde{\mathcal{E}}$ corresponds to a probabilistic function g such that $\mathbf{Pr}_x[f(x) \ne g(x)] \le O(\sqrt{\epsilon})$.

We further show that all of the above results extend to quantum channels $\mathcal{E}$ that are convex combinations of the channels $\{\mathcal{E}_i\}$ implemented by QAC^0 circuits: $\mathcal{E}(\rho) = \sum_i \alpha_i \mathcal{E}_i(\rho)$. Note that it is not necessarily true that $\mathcal{E}$ can be implemented by a QAC^0 circuit.

4 TECHNICAL OVERVIEW

The starting point for our results is the seminal work of Linial, Mansour, and Nisan [33] on the Fourier spectrum of constant-depth classical circuits; we refer the interested reader to the monographs [19, 38] for further background on the subject.

4.0.1 The Work of Linial, Mansour, and Nisan (LMN). Suppose $f : \{0,1\}^n \to \{0,1\}$ is a Boolean function computed by a depth-d AC^0 circuit of size s. Recall that f—viewed as a function $f : \{\pm1\}^n \to \{\pm1\}$—can be expressed as a real multilinear polynomial

$$f = \sum_{S\subseteq[n]} \widehat{f}(S)\chi_S \qquad \text{where} \qquad \chi_S(x) := \prod_{i\in S} x_i$$

which can be viewed as a Fourier expansion of f. The main technical result of [33], namely Lemma 7, is the following bound on the Fourier spectrum of f:

$$\sum_{|S|>k} \widehat{f}(S)^2 \le s\cdot 2^{-\Theta\left(k^{1/d}\right)} \qquad \text{for all } k\in[n], \tag{2}$$

which, intuitively, states that most "Fourier mass" of f lies on its low-degree part. Although this bound has been subsequently sharpened [7, 23, 24, 28, 48], for the sake of simplicity, this work will focus solely on the [33] bound given by Equation (2).

The primary technical ingredient used by [33] to prove the above bound is Håstad's celebrated switching lemma [22]. As an immediate consequence of Equation (2), one can obtain correlation bounds against parity as well as a learning algorithm (under the uniform distribution) for AC^0 circuits; we sketch both of these results below, see Section 4.5 of [38] for a thorough exposition.

Correlation Bounds Against Parity. Suppose $f : \{\pm1\}^n \to \{\pm1\}$ is computed by a depth-d size-s AC^0 circuit. Recall that $\mathrm{Parity}_n : \{\pm1\}^n \to \{\pm1\}$ is given by $\mathrm{Parity}_n := \chi_{[n]}$. As an immediate consequence of Equation (2),

$$\left|\widehat{f}([n])\right| \le s\cdot 2^{-\Theta\left(n^{1/d}\right)}.$$

Furthermore, using Proposition 1.9 from [38], is straightforward to check that

$$\Pr_{x\sim\{\pm1\}^n}\left[f(x) = \mathrm{Parity}_n(x)\right] \le \frac{1}{2} + s\cdot 2^{-\Theta\left(n^{1/d}\right)}.$$

So, if $d = O(1)$ and $s = \mathrm{poly}(n)$, then f can agree with the parity function on at most $\frac{1}{2} + \exp(-\Theta(n))$ fraction of inputs. Since guessing $\{0,1\}$ uniformly at random gives correlation $1/2$ with the parity function, this result implies that, with constant-depth circuits, one cannot do much better than random guessing.

Learning AC^0 Circuits. Taking $k = \Theta\left(\log^d\left(\frac{s}{\epsilon}\right)\right)$, Equation (2) implies that $\sum_{|S|>k}\widehat{f}(S)^2 \le \epsilon$. In other words, all but ϵ of f's "Fourier weight" lies on its low-degree coefficients. Based on this observation, [33] suggest a natural learning algorithm for AC^0 : Simply estimate all the low-degree Fourier coefficients $\{\widehat{f}(S)\}_{|S|\le k}$ to sufficiently high accuracy, and—writing $\widetilde{f}(S)$ for the estimate of $\widehat{f}(S)$—output the $\{\pm1\}$-valued function

$$\mathrm{sign}\left(\sum_{|S|\le k}\widetilde{f}(S)\chi_S\right). \tag{3}$$

This gives a quasipolynomial time algorithm for learning AC^0 circuits. In fact, it is known that, under a strong enough cryptographic assumption, quasipolynomial time is *required* to learn AC^0 circuits even with query access [30].[1]

4.0.2 Spectral Concentration of QAC^0 Circuits. Inspired by the classical importance of low-degree Fourier concentration, i.e. Equation (2), one might hope for an analogous notion of low-degree Pauli concentration in the quantum setting. As we saw earlier in the introduction via the example $U = X^{\otimes n}$, unitaries implemented by QAC^0 circuits do *not* have Pauli weight that is low-degree concentrated. Instead, we turn to analyzing the Pauli spectrum of QAC^0 *channels*.

[1] Note that the [33] learning algorithm only requires sample access to the function f, which is weaker than the class of algorithms with query access considered by Kharitonov [30].

The Pauli Decomposition of Channels. We define the *Pauli spectrum of channel* $\mathcal{E}$ as that of its Choi representation $\Phi_{\mathcal{E}}$:

$$\Phi_{\mathcal{E}} = \sum_P \widehat{\Phi_{\mathcal{E}}}(P)P.$$

To highlight that the Pauli spectrum of channels generalizes the classical Fourier spectrum of Boolean functions, note that when the channel computes a Boolean function $f : \{0,1\}^n \to \{0,1\}$, the Pauli spectrum of the Choi representation is closely related to the Fourier expansion of f:

$$\Phi_f = \frac{1}{2} I^{\otimes(n+1)} + \frac{1}{2} \sum_{S \subseteq [n]} \widehat{f}(S) Z_S \otimes Z \,. \tag{4}$$

Here, Z_S denotes $\bigotimes_{i=1}^n Z^{\mathbf{1}\{i \in S\}}$.

To further motivate our notion of Pauli spectrum, we return to our example $U = X^{\otimes n}$. Consider the channel

$$\mathcal{E}_U(\rho) = \mathrm{Tr}_{[n-1]}(U\rho U^\dagger)$$

that applies U to the input state and traces out all but the last qubit. As a check, it is readily verified that

$$\Phi_U = I^{\otimes n-1} \otimes \sum_{y,y' \in \{0,1\}} |y\rangle\langle y'| \otimes Z|y\rangle\langle y'|Z.$$

Thus $\mathcal{E}_U$ is "low-degree" as originally hoped for.

The Pauli Spectrum of QAC^0 *Channels.* Returning to quantum circuits, suppose $\mathcal{E}$ is implemented by a depth-d QAC^0 circuit acting on n input qubits and a auxiliary qubits. Writing s for the number of Toffoli gates acting in the circuit, this work proves a bound on the Pauli spectrum of $\Phi_{\mathcal{E}}$. Namely, for each $k \in [n+1]$, we show that

$$\sum_{|P|>k} |\Phi_{\mathcal{E}}(P)|^2 \leq O\left(s^2 2^{-k^{1/d}+a}\right). \tag{5}$$

Note the resemblance between Equation (5) and Equation (2) obtained by [33]. We now describe the proof of Equation (5).

For simplicity, we first describe the case when U is implemented by a circuit *without* any auxiliary qubits, i.e. when $a = 0$. In this case, the proof proceeds in two steps:

(1) We first establish that if a depth-d QAC^0 circuit does not have any Toffoli gates of width at least $k^{1/d}$, then it has no Pauli weight above level k; and
(2) Next, we show that deleting such "wide" Toffoli gates does not noticeably affect the action of the circuit.

A lengthy, albeit ultimately straightforward, calculation using standard analytic tools then establishes Equation (5) for the case when $a = 0$. We note that the proof of the first item above relies on a lightcone-type argument.

For the more general case where the circuit implementing U has a auxiliary qubits, we consider two cases, corresponding to *clean* auxiliary qubits (i.e. when the a qubits must start in the $|0^a\rangle$ state) and *dirty* auxiliary qubits (i.e. when there is no guarantee on the setting of the state of the a qubits). With clean auxiliary qubits, we can view the a auxiliary qubits as a part of the input to the circuit that we enforce to be set to $|0^a\rangle$ by postselecting the Choi representation; this results in the 2^a blow-up in Equation (5). In the dirty setting, however, we are able to incur no blowup; we defer discussion of this to the full version of the paper.

New Circuit Lower Bounds. As an immediate consequence of our spectral concentration bound on QAC^0, we obtain correlation bounds against parity and majority functions. This follows by:

(1) First, relating the classical Fourier expansion of these functions to the Pauli expansion of the Choi representations of channels implementing those functions (as in Equation (4)); and
(2) Then, showing that they cannot be approximated by QAC^0 with a small number of auxiliary qubits thanks to the spectral concentration as discussed above.

4.0.3 Learning QAC^0 *Circuits.* Our main learning result is an algorithm for learning channels with Pauli weight that is low-degree concentrated. Combining this with our low-degree concentration bound for QAC^0 channels (Equation (5)) immediately provides a learning algorithm for QAC^0 channels (with limited auxiliary qubits). Specifically, we show that quantum channels from n to 1 qubits that are implemented with a QAC^0 circuit and $o(\mathrm{poly}\log(n))$ auxiliary qubits can be learned[2] using quasipolynomial, i.e.

$$2^{O(\mathrm{poly}\log n)} \text{ samples.}$$

In comparison, naive tomography would require $2^{O(n)}$ samples.

A Quantum Low-Degree Algorithm. Our algorithm is inspired by—and closely follows the structure of—the classical low-degree algorithm introduced by [33].

The first step of the learning algorithm is to estimate all low-degree Pauli coefficients (i.e. all $P \in \mathcal{P}_n$ for $|P| \leq \mathrm{poly}\log(n)$) to sufficiently high accuracy. We consider two different learning models:

- **Choi state samples:** In the first model we are given copies of the Choi state $\rho_{\mathcal{E}} = \frac{\Phi_{\mathcal{E}}}{\mathrm{Tr}(\Phi_{\mathcal{E}})}$. We apply Classical Shadow Tomography [27] to learn the Pauli coefficients.
- **Measurement queries:** In this model, the learner is allowed to query an input state ρ and observable O and is given the measurement outcome of $\mathcal{E}(\rho)$ with respect to O. For this setting, we perform direct tomography on the quantum channel $\mathcal{E}$ for specially chosen input states ρ and measurement observables O.

While our approaches for both settings achieve similar sample complexity, the latter has the benefit of being more feasible (from an experimental implementation standpoint) than the former. Furthermore, the first approach can be thought of as analogous to learning from random labeled examples, and the second can be thought of as learning from query access to the function.

"Rounding" to CPTP maps. After obtaining estimates $\widetilde{\Phi}(P)$ for each of the Pauli coefficients $\widehat{\Phi}(P)$, the final step is to round $\widetilde{\Phi}_{\mathcal{E}} = \sum_P \widetilde{\Phi}_{\mathcal{E}}(P)P$ to a Choi representation of a valid quantum channel—that is, a completely-positive trace-preserving (CPTP) map. Since the set of all CPTP Choi representations is a convex set, we show that there exists a projection onto this set that will only reduce the distance of our learned state to the true Choi representation. However, it remains open whether there exists an algorithm to implement an exact rounding procedure in quasipolynomial time.

[2] By "learning a channel" we mean learning an approximation (in Frobenius distance) of its Choi representation.

Table 2: Lower bounds against QAC circuits. Here, "auxbits" refers to auxiliary qubits.

	Hard Function	Depth (d)	Restrictions	Guarantee
[16]	Parity_n	$O(1)$	$o(n)$ auxbits	Worst Case
[39]	Parity_n	2	None	Worst Case
[42]	Parity_n	2	None	Average Case
[42]	Parity_n	$O(1)$	$O(n/d)$ gates	Average Case
Theorem 2	Parity_n	$O(1)$	$\frac{1}{2}n^{1/d}$ auxbits	Average Case
Theorem 2	Majority_n	$O(1)$	$\frac{1}{2}n^{1/d}$ auxbits	Average Case

Note that this final step is analogous to the classical low-degree algorithm's use of the sign function, as in Equation (3).

5 RELATED WORK

Previous work on Pauli analysis. The original proposal for "quantum analysis of Boolean functions" came from [35], who proposed the study of Pauli decompositions (i.e. "Pauli analysis") of *Hermitian unitaries*. The key intuition is that these unitaries possess ± 1 eigenvalues, which can be interpreted as the outputs of a classical Boolean function. Recently, [43] extended seminal results from classical analysis of Boolean functions to this quantum setting. Furthermore, [13] extended these Pauli analysis techniques to the setting of quantum *non-Hermitian unitaries* for applications to quantum junta testing and learning. Unfortunately, however, as illustrated earlier in this section, the Pauli spectrum of a unitary U does not cleanly connect with the complexity of implementing U. In this work, we instead look to the Pauli decomposition of *channels*.

Recently, [3] also look to Pauli-analysis of quantum channels to extend the junta property testing and learning results of [13] to n-qubit to n-qubit channels. They employ Pauli-analytic techniques by proposing Fourier analysis of superoperators in the *Kraus representation*. Notably, their analysis is limited to n to n-qubit channels. In this work, we also consider Pauli analysis of superoperators. However, we instead propose the *Choi representations* of n-qubit to ℓ-qubit channels as our key objects of study. As discussed earlier in the section, this provides a definition of Pauli spectrum that connects more closely with computational complexity and allows us to prove spectral concentration, average-case lower bounds, and learning results for single-output-qubit QAC^0 circuits.

Previous work on QAC^0 *lower bounds.* Since Moore's paper [36] that originally defined the model, there has only been a smattering of lower bound results on QAC. We summarize known lower bounds against QAC^0 circuits below and in Table 2. We then compare them with our lower bound results.

Fang, et al. [16] established the first lower bounds on the QAC^0 model; in particular they proved that a depth-d QAC circuit cleanly computes the n-bit parity function with a auxiliary qubits, then $d \geq 2\log(n/(a+1))$. Here, "cleanly" means that the auxiliary qubits have to start and end in the zero state.

The key to their lower bound proof is a beautiful lemma (Lemma 4.2 of [16]): for all depth-d QAC circuits, there exists a subset S of $(a+1)2^{d/2}$ input qubits and a state $|\psi_S\rangle$ for that subset S, such that no matter what the other input qubits are set to, the output and auxiliary qubits result in the zero state. This immediately implies a lower bound for QAC circuits that cleanly compute the parity function: First, the clean computation property implies that without loss of generality the subset S is supported on non-auxiliary qubits. Second, if $d < 2\log(n/(a+1))$, then there exists a non-auxiliary input qubit i that is not fixed by $|\psi_S\rangle$, but the output qubit should depend on the state of the i'th qubit – except the output is already fixed to zero, a contradiction.

This lower bound is nontrivial as long as the number of auxiliary qubits is sublinear (i.e. $a = o(n)$), whereas our lower bound on the parity function can only handle up to $\sim n^{1/d}$ auxiliary qubits. On the other hand, the lower bound of [16] appears to be tailored to the setting where the circuit has to compute parity both exactly and cleanly. For a circuit that computes parity exactly (i.e. on all input strings), the clean computation property is without loss of generality because one can always save the output and then uncompute. When the circuit only computes parity approximately (e.g. on $\frac{1}{2}+\epsilon$ fraction of inputs), the clean computation property becomes an additional assumption.

Furthermore, the technique of [16] does not obviously extend to obtain average-case lower bounds: although there may be a fixing $|\psi_S\rangle$ of $(a+1)2^{d/2}$ input qubits that force the output qubit to be zero, such a fixing occurs with probability at most $2^{-(a+1)2^{d/2}}$ under the uniform distribution on the n input qubits – note that this is *exponentially* small in a and *doubly-exponentially* small in d. This directly implies that a depth-d QAC circuit with a auxiliary qubits cannot compute more than $1-2^{-(a+1)2^{d/2}}$ fraction of inputs. When $a = \omega(\log n)$ this fraction is extremely close to 1. By comparison our average case lower bound shows that QAC circuits with limited auxiliary qubits cannot compute parity on more than $\frac{1}{2}+2^{-\Omega(n^{1/d})}$ fraction of inputs.

We note[3] that the techniques of [16] also yield a lower bound on cleanly computing the majority function, as fixing a sublinear number of input bits is not enough to fix the majority function. However, for the same reasons as mentioned in the previous paragraph, it is unclear whether this argument can be extended to prove an average-case lower bound.

Later, Padé et al. [39] proved that no depth-2 quantum circuit (with *any* number of auxiliary qubits) can cleanly compute parity in the worst case. They prove this by carefully analyzing the structure of states that can be computed by depth-2 QAC circuits. Similarly it is unclear whether these techniques can be extended to the non-clean or approximate computation setting.

Rosenthal [42] proved that any average-case lower bound on QAC circuits (approximately) computing parity must use a bound on the number of auxiliary qubits; once there are *exponentially many* auxiliary qubits, then there is a depth-7 QAC circuit approximately computing parity. Furthermore, Rosenthal proved the following average-case lower bounds:

(1) A depth-d QAC circuit needs at least $\Omega(n/d)$ multiqubit Toffoli gates in order to achieve a $\frac{1}{2}+\exp(-o(n/d))$ approximation of parity, regardless of the number of auxiliary qubits.
(2) Depth-2 QAC circuits, with any number of auxiliary qubits, cannot achieve $\frac{1}{2}+\exp(-o(n))$ approximation of parity, even

[3] We thank an anonymous reviewer for pointing this out to us.

non-cleanly. This proves an average-case version of the lower bound of [39].

(3) A particular restricted subclass of QAC circuits (of which his depth-7 construction is an example) requires exponential size to compute parity, even approximately.

These are the first average-case lower bound results for QAC that we are aware of; however, they apply to restricted classes of QAC circuits and notably do not take into account the number of auxiliary qubits. As mentioned earlier, any general (average-case) lower bound on QAC circuits computing parity (for depths 7 and greater) must depend on the number of auxiliary qubits.

More recently, Slote [45] initiated the study of the closely related circuit class that is QNC^0 circuits followed by AC^0 post-processing, denoted $\mathrm{AC}^0 \circ \mathrm{QNC}^0$. Slote conjectures that polynomial-sized $\mathrm{AC}^0 \circ \mathrm{QNC}^0$ can not approximate parity, and shows that this is indeed the case when either the QNC^0 circuit has no auxiliary qubits, or when the AC^0 circuit has linear size. Perhaps surprisingly, the explicit connection between $\mathrm{AC}^0 \circ \mathrm{QNC}^0$ and QAC^0 is unclear: while QAC^0 circuits can certainly implement QNC^0 circuits, it is unknown whether they can implement AC^0 — it is, as far as we know, possible that QAC^0 is incomparable with both AC^0 and $\mathrm{AC}^0 \circ \mathrm{QNC}^0$. Nevertheless, for both QAC^0 and $\mathrm{AC}^0 \circ \mathrm{QNC}^0$, many existing techniques (the lightcone argument) fail for similar reasons. Slote's approach utilizes Fourier analysis of Boolean functions and draws connections to nonlocal games.

Related work on quantum learning. Efficient learning of quantum dynamics is a long-standing challenge in the field. Techniques such as quantum process tomography [34], which aim to fully characterize arbitrary quantum channels, require exponentially many data samples to guarantee a small error in the learned channel, for all possible channels.

One way of achieving sample-efficient quantum channel learning algorithms is performing full tomography on *specific classes* of quantum channels. By focusing on a specific class, rather than all possible quantum channels, there often exists nice structure which can be leveraged to reduce the number of data samples required to fully characterize channels in the class. For example, [3] showed that n-qubit to n-qubit quantum k-junta channels (acting non-trivially on at most k out of n qubits) can be learned to error ϵ, with high probability, via $O(4^k/\epsilon^2)$ samples. In our learning result, we focus on quantum channels with an arbitrary number of output qubits, and with "low-degree" Choi representations[4]. We show that a k-degree channel, involving $m = n + \ell$ total input and output qubits, can be learned to error ϵ, with high probability, via $\widetilde{O}((3m)^k/(4^\ell\epsilon))$ samples. Our result is incomparable to [3] since an n-qubit to n-qubit k-junta channel does not satisfy our notion of low-degree concentration. To our knowledge, this is the first work to analyze and offer a learning algorithm specific to channels with low-degree Choi representations. Furthermore, through our concentration result, we establish that QAC^0 circuits mapping to a single-qubit output ($\ell = 1$) lie in this circuit class, resulting in the first quasipolynomial learning algorithm for single-qubit-output QAC^0.

[4]We formally define the notion of a "low-degree" Choi representation in the full version of the paper.

An alternative approach for achieving sample-efficient learning algorithms is performing *partial* tomography on arbitrary quantum channels. For example, rather than full process tomography of a channel $\mathcal{E}$, [26] consider the task of learning the function $f(O_i, \rho) = \mathrm{Tr}(O_i\mathcal{E}(\rho))$ for a class of M observables $\{O_i\}_{i=1}^M$ and input state ρ. For an arbitrary quantum channel $\mathcal{E}$ and bounded-degree observables of spectral norm $\|O_i\| \leq 1$, they prove that $2^{O(\log(1/\epsilon)\log(n))}$ samples are sufficient to learn the function for all observables to error ϵ with high probability. At a high level, our result and that of [26] both establish and leverage Fourier concentration (i.e. low-degree approximation) to obtain efficient learning algorithms for quantum channels. However, our results operate in different settings. Namely, our work learns a low-degree approximation of the channel's Choi representation, whereas theirs learns low-degree approximations of the channel's Heisenberg-evolved observables $O_i^* = \mathcal{E}^\dagger(O_i)$, where $f(O_i, \rho) = \mathrm{Tr}(O_i^*\rho)$. [26] show that under a locally-flat input distribution, the Heisenberg-evolved observables of general channels are well approximated by low-degree observables. While this enables efficient learning of *any* quantum channel, restriction to locally-flat input distributions implies that, for quantum channels encoding classical Boolean functions, measurement expectations will be biased towards inputs which are not in the computational basis and, thus, uninformative. Our work instead obtains a sample-efficient learning result for the specific class of Choi representations of single-output-qubit QAC^0 circuits, with average-case guarantees according to the uniform distribution over computational basis states. To obtain this result, we prove low-degree concentration of QAC^0 Choi representations. This concentration result can further be shown to imply concentration of the channel's Heisenberg-evolved observables and, thus, could potentially be leveraged by the [26] procedure to offer a learning guarantee for single-qubit-output QAC^0 channels without restriction to locally-flat input distributions. It is an interesting direction of future research to formally relate the two works.

Finally, sample-efficient quantum channel learning algorithms can also be achieved by leveraging *quantum-enhanced* experiments. [12] proved exponential separations between learning algorithms with external quantum memory and those without. Building upon this, [9] recently demonstrated that full characterization of an unknown quantum channel's Pauli transfer matrix requires exponentially many channel queries in the case of classical processing and memories, but only polynomial samples in the case of quantum processing and memory. In this work, however, we do not consider quantum-enhanced experiments. Instead, we demonstrate that there exists a non-quantum-enhanced quasipolynomial learning algorithm for approximate characterization of the full Choi representation of single-output QAC^0 channels.

6 DISCUSSION

We believe that much remains to be discovered about the analytic properties of QAC^0 circuits. We list some natural concrete (and not-so-concrete) questions for future work below:

Improved Spectral Concentration? Arguably the most natural open question is to improve the dependence on the number of auxiliary qubits in Theorem 1 so as to get a lower bound against QAC^0 circuits with polynomially many auxiliary qubits. In fact, we conjecture the following improved spectral bound:

Conjecture 1 (Spectral concentration for QAC^0)**.** Suppose $\mathcal{E}$ is an n to 1-qubit quantum channel that is implemented by a depth-d QAC^0 circuit on n input qubits and $\mathrm{poly}(n)$ auxiliary qubits with s Toffoli gates. Then for all $k \in [n+1]$, we have

$$\sum_{P \in \mathcal{P}_{n+1} : |P| > k} \left|\widehat{\Phi_{\mathcal{E}}}(P)\right|^2 \leq \mathrm{poly}(s) \cdot 2^{-\Omega\left(k^{1/d}\right)}$$

In particular, we expect no dependence on the number of auxiliary qubits in our spectral bound. Note that Conjecture 1 would immediately imply an average-case lower bound for parity as well as a lower bound for majority against QAC^0 circuits with polynomially many auxiliary qubits, and extend the guarantees of our learning algorithm to this broader class of circuits. This would also match the classical bound on the Fourier spectrum of AC^0 circuits obtained by [33].

Improved Learning Algorithms? Another natural direction is to improve the runtime of our learning algorithm (see Theorem 3): while we obtain quasipolynomial sample complexity, we do not provide an explicit algorithm for the final Choi representation rounding step. We conjecture that there exists a quasipolynomial time algorithm implementing an exact rounding procedure, which would also achieve a quasipolynomial runtime for the procedure. Recall that the runtime of the [33] learning algorithm is (under a strong enough cryptographic assumption) known to be optimal [30].

Connections to State Synthesis Problems? Recent work of Rosenthal [42] relates the problem of computing parity to various state synthesis problems. Could analytic methods as employed in this paper be used to prove state (or unitary) synthesis lower bounds?

Connections to Pseudorandomness? Classically, circuit lower bounds have led to unconditional constructions of pseudorandom generators [37]. One could ambitiously hope for unconditional constructions of pseudorandom states against classes of shallow quantum circuits via circuit lower bounds.

An Emerging Analogy? This work adds to an emerging analogy between Fourier analysis in the classical setting of Boolean functions and "Pauli analysis" in the quantum setting of unitary operators or more generally quantum channels [3, 13, 35, 43, 50]. Given the tremendous success of Fourier analysis in classical complexity theory, we suspect that much remains to be discovered about the Pauli spectrum of quantum operations.

ACKNOWLEDGEMENTS

We thank Rocco Servedio for sagacious feedback. We thank Joseph Slote for discussions on connections between quantum circuits and Fourier analysis. We thank Daniel Grier and Gregory Rosenthal for helpful conversations. N.P. thanks Sergey Bravyi and Chinmay Nirkhe for thoughtful discussions. F.V. thanks Hsin-Yuan Huang for informative discussions on learning. We thank anonymous reviewers for their helpful feedback. We thank Mauricio Soler for helpful feedback and discussion. S.N. is supported by NSF grants IIS-1838154, CCF-2106429, CCF-2211238, CCF-1763970, and CCF-2107187. F.V. is supported by NSF grant DGE-2146752 and the Vannevar Bush Faculty Fellowship Program grant N00014-21-1-2941. N.P. and H.Y. are supported by AFOSR award FA9550-21-1-0040, NSF CAREER award CCF-2144219, and the Sloan Foundation. This work was partially completed while the authors were visiting the Simons Institute for the Theory of Computing.

REFERENCES

[1] Dorit Aharonov, Xun Gao, Zeph Landau, Yunchao Liu, and Umesh Vazirani. 2023. A Polynomial-Time Classical Algorithm for Noisy Random Circuit Sampling. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (Orlando, FL, USA) *(STOC 2023)*. Association for Computing Machinery, New York, NY, USA, 945–957. https://doi.org/10.1145/3564246.3585234

[2] Anurag Anshu, Nikolas P Breuckmann, and Chinmay Nirkhe. 2023. NLTS Hamiltonians from good quantum codes. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 1090–1096.

[3] Zongbo Bao and Penghui Yao. 2023. Nearly Optimal Algorithms for Testing and Learning Quantum Junta Channels. *arXiv preprint arXiv:2305.12097* (2023).

[4] Debajyoti Bera, Frederic Green, and Steven Homer. 2007. Small depth quantum circuits. *ACM SIGACT News* 38, 2 (2007), 35–50.

[5] Eric Blais. 2009. Testing Juntas Nearly Optimally. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing* (Bethesda, MD, USA) *(STOC '09)*. Association for Computing Machinery, New York, NY, USA, 151–158. https://doi.org/10.1145/1536414.1536437

[6] Dolev Bluvstein, Harry Levine, Giulia Semeghini, Tout T Wang, Sepehr Ebadi, Marcin Kalinowski, Alexander Keesling, Nishad Maskara, Hannes Pichler, Markus Greiner, et al. 2022. A quantum processor based on coherent transport of entangled atom arrays. *Nature* 604, 7906 (2022), 451–456.

[7] Ravi B Boppana. 1997. The average sensitivity of bounded-depth circuits. *Information processing letters* 63, 5 (1997), 257–261.

[8] M. Braverman. 2009. Poly-logarithmic independence fools AC^0 circuits. In *Proc. 24th Annual IEEE Conference on Computational Complexity (CCC)*. 3–8.

[9] Matthias C. Caro. 2023. Learning Quantum Processes and Hamiltonians via the Pauli Transfer Matrix. arXiv:2212.04471 [quant-ph]

[10] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. 2019. Pseudorandom generators from polarizing random walks. *Theory of Computing* 15, 1 (2019), 1–26.

[11] Eshan Chattopadhyay, Pooya Hatami, Shachar Lovett, and Avishay Tal. 2018. Pseudorandom generators from the second Fourier level and applications to AC0 with parity gates. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

[12] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. 2022. Exponential Separations Between Learning With and Without Quantum Memory. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. 574–585. https://doi.org/10.1109/FOCS52979.2021.00063

[13] Thomas Chen, Shivam Nadimpalli, and Henry Yuen. 2023. Testing and learning quantum juntas nearly optimally. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1163–1185. https://doi.org/10.1137/1.9781611977554.ch43

[14] Lior Eldar and Aram W Harrow. 2017. Local Hamiltonians whose ground states are hard to approximate. In *2017 IEEE 58th annual symposium on foundations of computer science (FOCS)*. IEEE, 427–438.

[15] Alexandros Eskenazis, Paata Ivanisvili, and Lauritz Streck. 2022. Low-degree learning and the metric entropy of polynomials. *arXiv preprint arXiv:2203.09659* (2022).

[16] Maosen Fang, Stephen Fenner, Frederic Green, Steven Homer, and Yong Zhang. 2003. Quantum lower bounds for fanout. *arXiv preprint quant-ph/0312208* (2003).

[17] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. 2002. Monotonicity testing over general poset domains. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*. 474–483.

[18] M. Furst, J. Saxe, and M. Sipser. 1984. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory* 17, 1 (1984), 13–27.

[19] Christophe Garban and Jeffrey E Steif. 2014. *Noise sensitivity of Boolean functions and percolation*. Vol. 5. Cambridge University Press.

[20] Pranav Gokhale, Samantha Koretsky, Shilin Huang, Swarnadeep Majumder, Andrew Drucker, Kenneth R Brown, and Frederic T Chong. 2021. Quantum fan-out: Circuit optimizations and technology modeling. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 276–290.

[21] Andrew Y Guo, Abhinav Deshpande, Su-Kuan Chu, Zachary Eldredge, Przemyslaw Bienias, Dhruv Devulapalli, Yuan Su, Andrew M Childs, and Alexey V Gorshkov. 2022. Implementing a fast unbounded quantum fanout gate using power-law interactions. *Physical Review Research* 4, 4 (2022), L042016.

[22] J. Håstad. 1986. *Computational Limitations for Small Depth Circuits*. MIT Press, Cambridge, MA.

[23] Johan Håstad. 2001. A slight sharpening of LMN. *J. Comput. System Sci.* 63, 3 (2001), 498–508.

[24] Johan Håstad. 2014. On the correlation of parity and small-depth circuits. *SIAM J. Comput.* 43, 5 (2014), 1699–1708.

[25] Peter Høyer and Robert Špalek. 2005. Quantum fan-out is powerful. *Theory of computing* 1, 1 (2005), 81–103.
[26] Hsin-Yuan Huang, Sitan Chen, and John Preskill. 2022. Learning to predict arbitrary quantum processes. *arXiv preprint arXiv:2210.14894* (2022).
[27] Hsin-Yuan Huang, Richard Kueng, and John Preskill. 2020. Predicting many properties of a quantum system from very few measurements. *Nature Physics* 16, 10 (Oct. 2020), 1050–1057. https://doi.org/10.1038/s41567-020-0932-7
[28] Russell Impagliazzo, William Matthews, and Ramamohan Paturi. 2012. A satisfiability algorithm for AC^0. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 961–972.
[29] A. Kalai, A. Klivans, Y. Mansour, and R. Servedio. 2005. Agnostically Learning Halfspaces. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science (FOCS)*. 11–20.
[30] Michael Kharitonov. 1993. Cryptographic hardness of distribution-specific learning. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*. 372–381.
[31] Subhash Khot, Dor Minzer, and Muli Safra. 2018. On Monotonicity Testing and Boolean Isoperimetric-type Theorems. *SIAM J. Comput.* 47, 6 (2018), 2238–2276.
[32] E. Kushilevitz and Y. Mansour. 1993. Learning decision trees using the Fourier spectrum. *SIAM J. on Computing* 22, 6 (1993), 1331–1348.
[33] N. Linial, Y. Mansour, and N. Nisan. 1993. Constant depth circuits, Fourier transform and learnability. *J. ACM* 40, 3 (1993), 607–620.
[34] M. Mohseni, A. T. Rezakhani, and D. A. Lidar. 2008. Quantum-process tomography: Resource analysis of different strategies. *Phys. Rev. A* 77 (Mar 2008), 032322. Issue 3. https://doi.org/10.1103/PhysRevA.77.032322
[35] Ashley Montanaro and Tobias J Osborne. 2008. Quantum boolean functions. *arXiv preprint arXiv:0810.2435* (2008).
[36] Cristopher Moore. 1999. Quantum circuits: Fanout, parity, and counting. *arXiv preprint quant-ph/9903046* (1999).
[37] Noam Nisan and Avi Wigderson. 1994. Hardness vs randomness. *Journal of computer and System Sciences* 49, 2 (1994), 149–167.
[38] Ryan O'Donnell. 2014. *Analysis of Boolean Functions*. Cambridge University Press. https://doi.org/10.1017/cbo9781139814782.013
[39] Daniel Padé, Stephen Fenner, Daniel Grier, and Thomas Thierauf. 2020. Depth-2 QAC circuits cannot simulate quantum parity. *arXiv preprint arXiv:2005.12169* (2020).
[40] Ran Raz and Avishay Tal. 2022. Oracle separation of BQP and PH. *ACM Journal of the ACM (JACM)* 69, 4 (2022), 1–21.
[41] Alexander A Razborov. 1987. 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 (1987), 333–338.
[42] Gregory Rosenthal. 2021. Bounds on the QAC^0 Complexity of Approximating Parity. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 185)*, James R. Lee (Ed.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 32:1–32:20. https://doi.org/10.4230/LIPIcs.ITCS.2021.32
[43] Cambyse Rouzé, Melchior Wirth, and Haonan Zhang. 2022. Quantum Talagrand, KKL and Friedgut's theorems and the learnability of quantum Boolean functions. *arXiv preprint arXiv:2209.07279* (2022).
[44] Kenneth Rudinger, Guilhem J Ribeill, Luke CG Govia, Matthew Ware, Erik Nielsen, Kevin Young, Thomas A Ohki, Robin Blume-Kohout, and Timothy Proctor. 2022. Characterizing midcircuit measurements on a superconducting qubit using gate set tomography. *Physical Review Applied* 17, 1 (2022), 014014.
[45] Joseph Slote. 2024. Parity vs. AC0 with simple quantum preprocessing. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.
[46] Joseph Slote, Alexander Volberg, and Haonan Zhang. 2023. Noncommutative Bohnenblust-Hille inequality in the Heisenberg-Weyl and Gell-Mann bases with applications to fast learning. *arXiv preprint arXiv:2301.01438* (2023).
[47] P. Smolensky. 1987. Information Processing in Dynamical Systems: Foundations of Harmony Theory. In *Parallel Distributed Processing: Volume 1: Foundations*, D. E. Rumelhart, J. L. McClelland, et al. (Eds.). MIT Press, Cambridge, 194–281.
[48] Avishay Tal. 2017. Tight bounds on the Fourier spectrum of AC^0. In *32nd Computational Complexity Conference (CCC 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
[49] Ruben Verresen, Nathanan Tantivasadakarn, and Ashvin Vishwanath. 2021. Efficiently preparing Schrödinger's cat, fractons and non-Abelian topological order in quantum devices. *arXiv preprint arXiv:2112.03061* (2021).
[50] Alexander Volberg and Haonan Zhang. 2023. Noncommutative Bohnenblust–Hille inequalities. *Math. Ann.* (2023), 1–20.
[51] Guoming Wang. 2011. Property testing of unitary operators. *Physical Review A* 84, 5 (2011), 052328.
[52] A. Yao. 1985. Separating the polynomial time hierarchy by oracles. In *Proceedings of the Twenty-Sixth Annual Symposium on Foundations of Computer Science*. 1–10.

Received 13-NOV-2023; accepted 2024-02-11