Random Unitaries in Constant (Quantum) Time

Ben Foxman
Yale University, New Haven, CT, USA

Natalie Parham
Columbia University, New York, NY, USA

Francisca Vasconcelos
UC Berkeley, CA, USA

Henry Yuen
Columbia University, New York, NY, USA

Abstract

Random unitaries are a central object of study in quantum information, with applications to quantum computation, quantum many-body physics, and quantum cryptography. Recent work has constructed unitary designs and pseudorandom unitaries (PRUs) using $\Theta(\log \log n)$-depth unitary circuits with two-qubit gates.

In this work, we show that unitary designs and PRUs can be efficiently constructed in several well-studied models of *constant-time* quantum computation (i.e., the time complexity on the quantum computer is independent of the system size). These models are constant-depth circuits augmented with certain nonlocal operations, such as (a) many-qubit `TOFFOLI` gates, (b) many-qubit `FANOUT` gates, or (c) mid-circuit measurements with classical feedforward control. Recent advances in quantum computing hardware suggest experimental feasibility of these models in the near future.

Our results demonstrate that unitary designs and PRUs can be constructed in much weaker circuit models than previously thought. Furthermore, our construction of PRUs in constant-depth with many-qubit `TOFFOLI` gates shows that, under cryptographic assumptions, there is no polynomial-time learning algorithm for the circuit class QAC^0. Finally, our results suggest a new approach towards proving that `PARITY` is not computable in QAC^0, a long-standing question in quantum complexity theory.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Quantum Information, Pseudorandomness, Circuit Complexity

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.61

Related Version *Full Version*: `https://arxiv.org/abs/2508.11487`

Funding NP is supported by the Google PhD fellowship. FV is supported by the Paul and Daisy Soros Fellowship for New Americans, NSF Graduate Research Fellowship Grant DGE2146752, and NSF Grant 2311733. HY is supported by AFOSR award FA9550-23-1-0363, NSF CAREER award CCF-2144219, NSF award CCF-232993, and the Sloan Foundation.

Acknowledgements BF thanks Dale Jacobs, Jackson Morris, Thomas Schuster, and Vishnu Iyer for helpful discussions, and HY for hosting during a visit to Columbia where this work was conducted. NP thanks Miranda Christ for helpful discussions about weak pseudorandom functions. HY thanks Harry Zhou and Soonwon Choi for helpful discussions.

1 Introduction

Unitary designs and pseudorandom unitaries (PRUs) are efficiently implementable ensembles of unitaries that mimic Haar-random unitaries; t-designs are *statistically* indistinguishable from the Haar measure up to t'th moments, and PRUs are *computationally* indistinguishable from Haar when queried by efficient quantum algorithms. Applications of unitary designs

and pseudorandom unitaries are diverse, ranging from quantum compiling [74], benchmarking [21, 51], quantum learning [38], quantum supremacy experiments [4], and quantum cryptography [31].

There is a long line of work focused on efficiently implementing unitary designs [10, 33, 34, 52, 45, 66] and PRUs [66, 50, 49]. Particularly striking are the "gluing-type" results of LaRacuente and Leditzky [45] and Schuster, Haferkamp, and Huang [66]. In these works, the authors show that several small, independent copies of t-designs or PRUs can be "glued" together, yielding random unitaries on larger numbers of qubits. This allows one to construct t-designs/PRUs in $O(\log n)$ depth on nearest-neighbor circuit geometries, and $O(\log \log n)$ depth in all-to-all circuits. Given that circuit depth corresponds to the time complexity of a computation, this results imply that (pseudo)random unitaries can be implemented in little time on a quantum computer.

Schuster et al. [66] also show that the Gluing Lemma is optimal: unitary designs and PRUs provably require $\Theta(\log \log n)$ depth in all-to-all architectures and $\Theta(\log n)$ depth in nearest-neighbor architectures. Given this, it may appear that the story is closed on minimizing the time complexity of unitary designs and PRUs. However, this is only the limit when we consider *unitary* circuits composed of single- and two-qubit gates. In this paper we show that t-designs and PRUs can be implemented in *constant time* in several well-motivated models of quantum computation. We then describe several implications of our results, ranging from possible experimental implementations, quantum complexity theory, and limits on quantum learning.

Models of constant-time computation

Constant-depth quantum circuits are a model of *constant-time* quantum computation – the time complexity on the quantum computer is independent of the system size). For state and unitary synthesis tasks, constant-depth circuits with only local unitary gates are quite limited in power; for example, they cannot prepare long-range entangled states (like the GHZ state). However, there has been significant interest in understanding the power of constant-depth quantum circuits augmented with complex operations such as many-qubit gates or intermediate measurements with feedforward control. The addition of nonlocal operations appears to unlock significant computational advantages for constant-depth circuits. Some models include:

1. Short-time Hamiltonian evolutions. Such operations are natural in analog quantum simulators [59], and are capable of entangling many qubits simultaneously [26].
2. Constant-depth circuits with mid-circuit measurements (i.e., intermediate measurements, followed by classical postprocessing of measurement outcomes to control subsequent layers of gates). This model is important for quantum error-correction [67], and has been shown to be capable of generating exotic quantum states with long-range entanglement [72, 12]. Here we treat classical computation as a free operation (i.e., costing no time), which is a reasonable assumption in many cases.
3. Constant-depth circuits with arbitrary single-qubit gates and many-qubit `FANOUT` gates (which maps $|s, x_1, \ldots, x_k\rangle$ to $|s, s \oplus x_1, \ldots, s \oplus x_k\rangle$). This gives rise to the so-called QAC^0_f circuit class, which has been studied in quantum complexity theory as a quantum analogue of the classical circuit class AC^0 [53]. The addition of the `FANOUT` gate makes constant-depth circuits surprisingly powerful; for example, the `PARITY` function and even the quantum Fourier transform can be computed with QAC^0_f circuits [37].
4. Circuits with constant T*-depth*, that is, arbitrary Clifford computation interleaved with a constant number of layers of T gates. Such circuits have been studied in the context of fault-tolerant quantum computation [27], nonlocal quantum computation [70], and quantum

circuit lower bounds [58]. This is a reasonable model of constant-time computation in the fault-tolerance setting, because the cost of performing logical T gates dominates that of Clifford operations.

These models might appear very different from each other, but in fact some of them are closely related. For example, constant T-depth circuits, when given access to a catalytic resource state, can simulate QAC^0_f [43], which in turn is equivalent to constant-depth circuits with mid-circuit measurements where the classical processing consists of parity computations [12].

The equivalence of these models greatly extends the flexibility of quantum algorithm design. For example, one could design an efficient quantum algorithm in the QAC^0_f model (which may be conceptually easier to think about), and then immediately obtain an efficient implementation in the mid-circuit measurement model, which is experimentally realizable [18]. We take this approach to show that unitary designs and pseudorandom unitaries can be implemented in constant (quantum[1]) time.

1.1 Our Contributions

First we show that unitary designs and PRUs can be implemented by polynomial-size QAC^0_f circuits. Here, we consider two flavors of designs and PRUs, *standard* and *strong*. Standard means that the unitaries appear Haar-random to algorithms that can make (adaptive) queries to the forwards direction of the unitary, whereas strong means that the unitaries appear Haar-random even when the algorithm can make queries to the inverse unitary as well [50].

This theorem circumvents the lower bounds of [66] which only apply to circuits with single- and two-qubit unitary gates.

▶ **Theorem 1.1** (Random Unitaries in QAC^0_f). *Let n denote the number of input qubits. The following can be implemented in QAC^0_f, the class of constant-depth circuits with arbitrary single-qubit gates and many-qubit* `FANOUT` *gates:*

- *Standard $1/\operatorname{poly}(n)$-approximate t-designs can be implemented with $n \operatorname{poly}\log n \cdot \operatorname{poly}(t)$-size circuits; and*
- *Strong $\exp(-\Omega(n))$-approximate t-designs can be implemented with $\operatorname{poly}(n, t)$-size circuits.*

Furthermore, assuming subexponential post-quantum security of the Learning With Errors (LWE) problem, the following can also be implemented in QAC^0_f:

- *Standard PRUs can be implemented with $n \operatorname{poly}\log n$-size circuits; and*
- *Strong PRUs can be implemented with $\operatorname{poly}(n)$-size circuits.*

The precise definitions of approximate unitary designs and PRUs can be found in Section 2.3, and we prove Theorem 1.1 in Section 3. We stress that while the designs are implemented by constant-depth QAC^0_f circuits, the pseudorandomness holds against *all* algorithms that make at most t queries; in particular the algorithms can perform arbitrarily deep computation in between queries.

The result for t-designs and PRUs are stated and proved in Theorem 3.4 and Theorem 3.7 respectively. We note that a similar approach was used by Chia, Liang, and Song [16] to obtain random unitaries in QAC^0_f, under slightly different assumptions.

The close connection between QAC^0_f and circuits with mid-circuit measurements yields the following corollary: unitary designs and PRUs can also be constructed in constant-depth 2D-local circuit geometries using nearest-neighbor gates, intermediate measurements, and feedforward operations.

[1] The parenthetical is meant to emphasize that the *quantum* part of the circuit has constant depth, and we treat the classical processing, which will be simple parity computations, as free.

▶ **Corollary 1.2** (Random unitaries in the intermediate measurement model)**.** *Approximate t-designs and PRUs can be implemented by constant-depth quantum circuits with nearest-neighbor two-qubit gates on a 2D lattice, intermediate measurements, and classical feedforward consisting of computing parities of measurement outcomes. The circuits for t-designs have size $n \operatorname{poly}\log n \cdot \operatorname{poly}(t)$, and the circuits for PRUs have size $n \operatorname{poly}\log(n)$. In the same model, strong t-designs and PRUs can be implemented by circuits with size $\operatorname{poly}(n, t)$ and $\operatorname{poly}(n)$ respectively.*

We prove Corollary 1.2 in Section 3.4.

▶ Remark 1.3. It is important to note that, even though the mid-circuit measurements yield random outcomes, Corollary 1.2 gives a *deterministic* implementation of every unitary in the t-design/PRU ensemble. The measurement outcomes are only used for simple Pauli corrections, but don't change the overall unitary that is being performed.

Theorem 1.1 and Corollary 1.2 provides another illustration of the power of models of constant-time quantum computation. There is a long line of work demonstrating the various capabilities of QAC^0_f circuits and circuits with mid-circuit measurement, from the quantum Fourier transform [37] to the preparation of novel states from condensed matter physics [72, 12]. We can now add t-designs and PRUs to this list.

Furthermore, our constructions suggest the possibility of highly efficient experimental implementations of random unitaries. Random unitaries have been extremely useful in a variety of experimental settings, such as benchmarking [21, 51], demonstrations of quantum supremacy [4], and explorations of quantum chaos [7] and quantum thermalization [29]. A number of experimental platforms are starting to support mid-circuit measurements and feedforward control [18, 40], which may bring the protocol of Corollary 1.2 within the realm of near-term implementation.

Random unitaries in QAC^0

In [53], Moore also introduced a variant of QAC^0_f called QAC^0 – note the lack of the f subscript. This model consists of constant-depth circuits with arbitrary single-qubit gates and (instead of the `FANOUT` gate) many-qubit `TOFFOLI` gates, which maps $|x_1, \ldots, x_k, s\rangle$ to $|x_1, \ldots, x_k, s \oplus \texttt{AND}(x_1, \ldots, x_k)\rangle$. Both QAC^0 and QAC^0_f were defined as natural analogues of the classical circuit class AC^0, an extremely well-studied circuit model in theoretical computer science [28].

QAC^0 can be efficiently simulated by QAC^0_f circuits[2], but it is believed that QAC^0 is much weaker than QAC^0_f. It is conjectured, for example, that the `PARITY` function cannot be computed by polynomial-size QAC^0 circuits, implying that $\mathsf{QAC}^0 \neq \mathsf{QAC}^0_f$. Despite QAC^0 appearing to be such a weak circuit class, we show that such circuits can nevertheless efficiently implement t-designs for $t = O(1)$ and PRUs with inverse-polynomial approximation error.

▶ **Theorem 1.4** (Random unitaries in QAC^0)**.** *Let n denote the number of input qubits. Let $\delta > 0$ and let $k \in \mathbb{N}$ be constants. The following can be implemented in QAC^0, the class of constant-depth circuits with arbitrary single-qubit gates and many-qubit `TOFFOLI` gates:*

[2] Moore originally defined QAC^0_f to include *both* `FANOUT` and `TOFFOLI` gates, which meant that QAC^0 is contained in QAC^0_f by definition. However, Takahashi and Tani later showed that `TOFFOLI` gates can be simulated using `FANOUT` gates [71].

- *Standard n^{-k}-approximate t-designs can be implemented with depth $O((\log k)/\delta)$, size $O(n^{1+\delta t})$ circuits.*

Furthermore, assuming subexponential post-quantum security of the Learning With Errors (LWE) problem, the following can also be implemented in QAC^0*:*

- *Standard PRUs can be implemented up to n^{-k} error, by depth $O((\log k)/\delta)$, size $O(n^{1+\delta})$ circuits.*

▶ Remark 1.5. We note that the PRUs in Theorem 1.4 can only be implemented by QAC^0 circuits up to inverse polynomial error (whereas the QAC^0_f circuits in Theorem 1.1 can implement PRUs with *exponentially* small error). This makes the PRU construction less relevant for cryptographic applications, perhaps, but we believe it still has the following use cases:

1. The QAC^0 t-design constructions only look random up to $t = O(1)$ moments; however the PRU constructions look random (to computationally-bounded adversaries) up to $O(n^k)$ queries, and
2. The PRU constructions imply lower bounds on the complexity of *learning* QAC^0 circuits, which we elaborate on shortly.

We find Theorem 1.4 interesting for several reasons. The first is that there seems to be no classical analogue. Pseudorandom functions cannot be efficiently implemented by classical AC^0 circuits [47], which are constant-depth circuits with unbounded fan-in `AND` and `NOT` gates (i.e., the classical analogue of QAC^0). Thus *quantum* pseudorandomness appears easier to achieve in a weak quantum circuit model than *classical* pseudorandomness in an analogous classical circuit model. We discuss in more detail in Section 1.2.

A second reason is that, to the best of our knowledge, Theorem 1.4 demonstrates the first example of an interesting task implementable in QAC^0. Moreover, for constant t-designs and PRUs, the number of ancillae is $n^{1+1/O(d)}$, where d is the depth of the QAC^0 circuit. By a result of Anshu et al. [1], the `PARITY` function cannot be computed by QAC^0 circuits with $n^{1+1/\exp(d)}$ ancillae. Therefore, there number of ancillae in our constructions are quite close to the current best-known `PARITY` lower bounds.

A third reason is due to a connection with quantum complexity theory. Given Theorem 1.4, one may naturally wonder about whether QAC^0 can implement strong designs or strong PRUs? Fascinatingly, this connects to a long-standing open question in quantum circuit complexity: can QAC^0 efficiently simulate QAC^0_f, or equivalently, is the `PARITY` function computable in QAC^0 [53]? If $\mathsf{QAC}^0 = \mathsf{QAC}^0_f$, then by Theorem 1.1 certainly strong designs/PRUs can be implemented in QAC^0. Conversely, if one can show that strong designs or PRUs *cannot* be implemented by QAC^0 circuits, then this proves $\texttt{PARITY} \notin \mathsf{QAC}^0$.

Hardness of learning QAC^0

A central problem in quantum learning theory is reconstructing the description of a quantum circuit given only black-box query access. As alluded to above, Theorem 1.4 has direct implications for learning QAC^0 circuits – a topic explored in several prior works [54, 6, 73]. In particular, assuming the hardness of the Learning with Errors (LWE) problem, the following theorem leverages our QAC^0 PRUs to establish super-polynomial time lower bounds for *average-case* learning of polynomial-size QAC^0 circuits.

▶ **Theorem 1.6** (QAC^0 Learning Lower-Bound)**.** *Assuming subexponential post-quantum security of LWE, $n^{\omega(1)}$ time is necessary to learn QAC^0 unitaries with $O(n^\delta)$ ancillae to average-case distance*[3]*, for any constant $\delta > 0$.*

[3] Here average-case distance refers to the error in average gate-fidelity according to Haar random inputs, as described in depth in [73].

We emphasize that this is the first *average-case* learning lower-bound for QAC^0, with previous work by [39, 73] establishing an exponential sample worst-case learning lower-bound.

Construction overview

We provide a brief overview of our approach. To construct unitary designs and PRUs in QAC^0_f (Theorem 1.1), we parallelize a variant of the so-called PFC construction of random unitaries, first introduced by [52] and further studied by [50, 20]. Here, P, F and C represent a random permutation matrix, a random diagonal matrix, and a random Clifford operator respectively. By combining these ensembles together, one obtains both t-designs and PRUs. Furthermore, the inclusion of a further random Clifford (the $CPFC$ ensemble) produces *strong* t-designs and *strong* PRUs [50]. We show that each of the P, F, and C ensembles can be implemented in QAC^0_f; this exploits the surprising power of the `FANOUT` gate [37, 30].

To obtain random unitaries in QAC^0 (without the `FANOUT` gate), we modify Morris and Grier's *exponential*-size QAC^0 implementation of the `FANOUT` gate [30] to convert the QAC^0_f constructions above into exponential-size QAC^0 circuits. We then scale these down to logarithmically-many input qubits (yielding polynomial-size QAC^0 circuits), and then apply the "Gluing Lemma" of [66] to recover a unitary design on n qubits. Random unitaries constructed via the Gluing Lemma are not secure when inverse queries are allowed, so Theorem 1.4 only yields the standard notion of designs (rather than strong designs).

1.2 Discussion

We now elaborate on the implications and consequences of our constructions of random unitaries in QAC^0 and QAC^0_f.

Classical versus Quantum Pseudorandomness

The quantum circuit models QAC^0 and QAC^0_f were first introduced by Moore to investigate quantum analogues of the classical circuit class AC^0, constant-depth classical circuits with unbounded fan-in `AND` and `NOT` gates. The study of AC^0 has been extremely fruitful in theoretical computer science, yielding celebrated lower bounds such as $\texttt{PARITY} \notin \mathsf{AC}^0$ [28, 35], powerful techniques such as random restrictions and the polynomial method [35, 61], and subsequent developments in learning theory and complexity theory that continue to this day [47, 11, 60].

The existence of random unitaries in QAC^0 (Theorem 1.4) is particularly surprising given that an analogous result provably does not hold for AC^0 [47]. Perhaps one interpretation is this: quantum `AND` (i.e., `TOFFOLI`) gates appear more effective at generating *quantum randomness* than classical `AND` gates are at generating *classical randomness*.

Towards Quantum Circuit Lower Bounds

Is the `PARITY` function computable in QAC^0? The question remains stubbornly open 25 years after Moore initially posed the question. The classical techniques for proving $\texttt{PARITY} \notin \mathsf{AC}^0$ do not seem to carry over to the quantum setting. Lower bounds have been given in some restricted settings, such as when the circuit has depth at most 2 [57, 64, 25], or the number of ancillae is limited [24, 54, 1]. However, even depth 3 lower bounds remain open.

In this work, we show that $\texttt{PARITY} \notin \mathsf{QAC}^0$ (both the unitary *and* the decision problem) is implied by the *non-existence* of strong random unitaries in QAC^0. While a formally a harder task, the connection with random unitaries may unlock new approaches to proving QAC^0 lower bounds. We elaborate on these approaches in the next section.

Table 1 Estimating properties with forward queries to just U, or queries to both U and $U^\dagger$. In the latter case, superpolynomial queries are required for each task when only forward queries are allowed.

Access Model	Capability
U only	Collision probability, purity of subsystems, average fidelity across inputs, overlap with Pauli operators [38]
U and $U^\dagger$	Determining lightcones [66], entanglement entropy [32, 14], OTOCs [19], displacement amplitudes [44]

Implications of Inverse Access

Our constructions of random unitaries in QAC^0 are secure against adversaries with only forward access to U. When `FANOUT` gates are allowed, we obtain random unitaries which are secure even with inverse access. The existence of forward-secure (pseudo)random unitaries in quasipolynomial size QAC^0 implies that all statistics in the first row of Table 1 must approximately agree with the Haar-random values. The same is true of random unitaries in QAC^0_f with respect to *all* properties in the table. Therefore, one cannot hope to prove lower bounds against QAC^0 or QAC^0_f by showing that these properties significantly differ from the expected Haar-random values. This is reminiscent of the *natural proofs barrier* introduced by Razborov and Rudich [62], which has been considered in other contexts in quantum information [22, 15].

Learning Lower-Bounds

Classically, [47] leveraged low-degree Fourier concentration to derive the first efficient (quasipolynomial) sample and time complexity algorithm for learning AC^0 functions. Assuming the classical hardness of factoring, [42] proved that quasipolynomial sample complexity of learning AC^0 is optimal.

On the quantum side, there has been substantial recent progress on learning QAC^0 [54, 73] gave a quasipolynomial time algorithms for learning outputs of QAC^0 circuit with few ancillae. Furthermore, assuming a quantum analog of the LMN Theorem, they showed that the algorithm can be extended to any polynomial-size QAC^0 circuit. Assuming subexponential-security of LWE, we show that quasipolynomial time complexity is near optimal for average-case learning of polynomial-size QAC^0 circuits.

1.3 Conclusion and Further Directions

Overall, our results highlight the surprising power of constant-depth circuits with many-qubit gates, suggesting the shallow quantum circuits may be even more powerful than previously thought. Since many quantum technologies are expected to support many-qubit gates [69, 9, 75], our results suggest that a staple of quantum information processing, (pseudo)random unitaries, may be realizable in constant time on future quantum hardware.

We believe that our results suggest some interesting directions for future investigation, which we list below.

Inverse Security and QAC^0

As mentioned in Section 1.2, proving that strong t-designs are not in QAC^0 is a path to proving $\texttt{PARITY} \notin \mathsf{QAC}^0$. In contrast with recent approaches [54, 1] which aim to generalize classical techniques to the quantum setting, this route would not require such an analogy.

Could properties in Table 1 be helpful for ruling out strong t-designs in QAC^0, and by extension proving $\texttt{PARITY} \notin \mathsf{QAC}^0$?

Weakening Assumptions for Random Unitaries

Even though AC^0 does not contain PRFs, there are candidate constructions of *weak* PRFs in AC^0 [8], which are secure against adversaries who only receive random $(x, f(x))$ samples as opposed to arbitrary query access. Could PRUs be built from post-quantum *weak* PRFs? Currently, it is only known how to build PRUs from standard quantum-secure PRFs, which are secure against adversaries making adaptive queries.

The "Quantum Natural Proofs Barrier"

Table 1 lists many properties which, by our results, must agree with Haar-random values for QAC^0 and QAC^0_f. This rules out several approaches for proving lower bounds against these classes. Are there other approaches for proving quantum circuit lower bounds which circumvent these barriers?

How Easy is Generating Quantum Pseudorandomness?

A broader takeaway from our paper is the idea that random unitaries can be implemented in a much wider range of constant-time models of quantum computation than previously thought. Could current-day programmable analog quantum simulators, which can perform globally entangling Hamiltonian evolutions [59], be sufficient to generate quantum pseudorandomness in constant time? Recent developments in using optical cavities with neutral atoms may offer further avenues to perform strongly coupled interactions between large numbers of qubits [46, 41] – could such interactions be used in place of the $\texttt{TOFFOLI}$ or $\texttt{FANOUT}$ gates to efficiently implement designs and PRUs in neutral atom arrays? Theoretically, can we characterize the weakest possible entangling operation that suffices to generate quantum pseudorandomness in constant time? These are interesting questions that we leave for future work.

2 Background

In this paper, exp and log denote base-2 exponent and logarithm. $\mathrm{poly}(n) = \bigcup_{c\in\mathbb{N}} O(n^c)$. A function $f(n)$ is *negligible* if $f(n) \in o(1/p(n))$ for any polynomial p. A function is *quasipolynomial* if it is of the form $n^{\mathrm{poly}\log n}$.

2.1 Quantum Computation

We assume basic familiarity with the formalism of quantum computing, including states, unitaries, and measurements. In this section, we briefly recall concepts that will be important for this work. First, the *Pauli string* on n qubit unitaries consists of all n-fold tensor products of the single-qubit matrices $\{I, X, Y, Z\}$. The set of Pauli strings is an orthonormal basis for the vector space of $2^n \times 2^n$ matrices, under the Hilbert-Schmidt inner product $\langle A, B\rangle = \frac{1}{2^n}\mathrm{Tr}(A^\dagger B)$. A *Clifford operator* is a unitary that normalizes the Pauli group, i.e. $CPC^\dagger$ is in the Pauli group if P is. The *trace distance* between two quantum states ρ and σ is defined as $\|\rho - \sigma\|_{td} := \frac{1}{2}\|\rho - \sigma\|_1$ where $\|A\|_1 = \mathrm{Tr}[\sqrt{A^\dagger A}]$ denotes the trace norm. The *Fuchs-van-de Graaf inequality* relates the trace distance $\|\rho - \sigma\|_{td}$ and the *fidelity* $F(\rho, \sigma) := (\mathrm{Tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})^2$ via the bound $\|\rho - \sigma\|_{td} \leq \sqrt{1 - F(\rho, \sigma)}$ If $\rho = |\psi_1\rangle\langle\psi_1|$ and $\sigma = |\psi_2\rangle\langle\psi_2|$, then $F(\rho, \sigma) = |\langle\psi_1|\psi_2\rangle|^2$.

Given a bipartite state ρ_{AB} on $\mathcal{H}_A \otimes \mathcal{H}_B$, the *partial trace* over subsystem B is the map $\mathrm{Tr}_B : \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B) \to \mathcal{L}(\mathcal{H}_A)$ defined by $\mathrm{Tr}_B(\rho_{AB}) = \sum_j (\mathbb{I}_A \otimes \langle j|)\rho_{AB}(\mathbb{I}_A \otimes |j\rangle)$ where $\{|j\rangle\}$ is an orthonormal basis for $\mathcal{H}_B$. The result is the *reduced density matrix* on $\mathcal{H}_A$, describing the marginal state when subsystem B is ignored.

Finally, the *Haar measure*, denoted μ_{2^n}, is the unique left and right unitarily invariant probability measure over the unitary group $\mathcal{U}(2^n)$.

2.2 Quantum Circuits

We assume familiarity with the circuit model of quantum computation, see [55] for an introduction. Quantum circuits can be used to solve classical decision problems, and quantum state or unitary synthesis problems. In this work, we consider both types of computation:

▶ **Definition 2.1.** *A quantum circuit C computes a function $f : \{0,1\}^n \mapsto \{0,1\}^m$ with average-case error ε if $\Pr_{x\in\{0,1\}^n}[C(x) = f(x)] \geq 1 - \varepsilon$. Here, $C(x)$ is a random variable denoting the measurement of m designated output qubits in the standard basis. The randomness is taken over both the input $|x\rangle$ and the randomness in the measurement.*

▶ **Definition 2.2.** *An $n + a$ qubit circuit C ε-implements an n-qubit unitary U if*

$$\left|\langle \psi, 0^a| C^\dagger (U|\psi\rangle \otimes |0^a\rangle)\right|^2 \geq 1 - \varepsilon$$

for every n-qubit state $|\psi\rangle$.

When $\varepsilon = 0$, we simply say C implements U. Approximate implementations have a composition property: we defer the proof to Lemma A.1.

▶ **Lemma 2.3.** *Assume that C ε_0-implements U and D ε_1-implements V. Then, CD $(\varepsilon_0 + \varepsilon_1)$-implements UV.*

In this work, we study quantum analogs of the classical circuit class AC^0, which are constant-depth circuits composed of unbounded fan-in AND/OR gates and NOT gates. Because these circuits are classical, bits can be freely copied and reused throughout the computation, and a single bit may serve as input to multiple gates in the next layer. However, since quantum circuits are unitary, any output qubit can be used in at most one subsequent gate. Moreover, due to the no-cloning theorem, quantum information cannot be freely copied. As a result, implementing a "quantum FANOUT"[4] requires explicitly including a FANOUT gate in the gate set. This distinction is captured by two different generalizations of AC^0 to the quantum setting: QAC^0 and QAC^0_f.

▶ **Definition 2.4** (QAC^0)**.** *A QAC^0 circuit is a constant-depth circuit consisting of arbitrary single-qubit gates and* TOFFOLI *gates, where*

$$\texttt{TOFFOLI}\,|x_1, \ldots, x_k, t\rangle = |x_1, \ldots, x_k, t \oplus \texttt{AND}(x_1, \ldots, x_k)\rangle\,.$$

Importantly, we allow the number of "control" qubits k to be arbitrary.

Unless otherwise specified, QAC^0 circuits have polynomial size, which is sometimes taken to be the definition of the class[5].

[4] Quantum information can only be copied in a fixed basis due to the no-cloning theorem.

[5] Formally, we consider QAC circuit families $\{C_n\}$ indexed by input size n, each of which has $\mathrm{poly}(n)$ size and $O(1)$ depth.

▶ **Definition 2.5** (QAC^0_f). *A QAC^0_f circuit is a constant depth circuit consisting of arbitrary single-qubit gates,* TOFFOLI *gates, and* FANOUT *gates, where*

$$\texttt{FANOUT}\,|s, x_1, \ldots, x_k\rangle = |s, s \oplus x_1, \ldots, s \oplus x_k\rangle\,.$$

Like QAC^0, QAC^0_f circuits have polynomial size unless specified otherwise. Since FANOUT allows classical bits to be copied, QAC^0_f can simulate polynomial size, constant-depth AC circuits, i.e. AC^0. However, unlike AC^0, QAC^0_f circuits can *exactly* compute the PARITY function

$$\texttt{PARITY}\,|x_1, \ldots, x_k, t\rangle \mapsto |x_1, \ldots, x_k, t \oplus \bigoplus_{i=1}^{k} x_i\rangle$$

due to the following identity:

▶ **Fact 2.6.** $H^{\otimes n}\texttt{PARITY}_{A_1\ldots,A_{n-1},B}H^{\otimes n} = \textit{FANOUT}_{B,A_{n-1},\ldots,A_1}$.

Building off of this identity, a long line of work [17, 71, 12, 30, 68] has shown that many other important operations can be simulated exactly by constant-depth circuits with FANOUT gates. Of particular interest to us is the $\texttt{THRESHOLD}_t$ gate, which flips a target bit if and only if the Hamming weight of the input is at least t. Amazingly, constant-depth THRESHOLD circuits (with arbitrary values of t) can be simulated exactly by QAC^0_f circuits [71]. We can state this formally by introducing the complexity class TC^0:

▶ **Definition 2.7** (TC^0). TC^0 *is the class of polynomial size, constant depth classical circuits with unbounded* AND, OR, *and* $\textit{THRESHOLD}_t$ *gates for arbitrary t.*

▶ **Theorem 2.8** ([12, 71]). *For any $f \in \mathsf{TC}^0$, the unitary operator $U_f : |x\rangle \mapsto (-1)^{f(x)}|x\rangle$ can be implemented by a QAC^0_f circuit.*

Theorem 2.8 is particularly interesting in the context of random unitaries, since constant-depth circuits with threshold gates can implement candidate quantum-secure one way functions.

2.3 Random Unitaries

Haar-random unitary operators play a central role in quantum information [4, 66, 56, 65]. In order to efficiently implement random unitaries, two main relaxations are typically considered: unitary t-designs [21] and pseudorandom unitaries [23]. The former restricts the number of times the algorithm queries U, while the latter restricts the algorithm's runtime.

▶ **Definition 2.9** ((Strong) Approximate Unitary t-designs). *A family of unitaries $\mathcal{U} = \{U_{n,k}\}_{n,k}$ is a* (strong) ε-approximate unitary t-design *if for all n and for any quantum algorithm Q making at most t queries to U (and its inverse $U^\dagger$),*

$$\left| \Pr_{k\sim\mathcal{K}}[Q^{U_{n,k},\,(U^\dagger_{n,k})}(1^n) = 1] - \Pr_{U\sim\mu_N}[Q^{U,\,(U^\dagger)}(1^n) = 1] \right| \leq \varepsilon$$

In the literature, this definition coincides with the recently defined notion of *measurable*-error approximate t-design [20], to distinguish it from the weaker notion of an *additive* error approximate t-design (which is secure only against adversaries making parallel queries) and the stronger notion of *relative* error (which is indistinguishable from Haar-random even for unphysical measurements). We will not need these distinctions, so we do not disambiguate.

▶ **Definition 2.10** ((Strong) Pseudorandom Random Unitaries). *A family of unitaries* $\mathcal{U} = \{U_{n,k}\}_{n,k}$ *is* (strongly) ε-secure against t-time quantum adversaries *if for all* n *and for any quantum algorithm* Q *running in time* t,

$$\left| \Pr_{k\sim\mathcal{K}}[Q^{U_{n,k},(U_{n,k}^\dagger)}(1^n) = 1] - \Pr_{U\sim\mu_N}[Q^{U,(U^\dagger)}(1^n) = 1] \right| \leq \varepsilon$$

When ε *is a negligible in* n *and* t *is superpolynomial in* n, *we say that* $\mathcal{U}$ *is a (strong)* pseudorandom unitary *(PRU).*

We take care to distinguish the error of a t-design or PRU, from the error incurred by a circuit implementation of such unitaries. For example, we will refer to "ε_0-approximate t-designs, which are ε_1-implemented by a circuit C".

In this paper, we make use of a particular PRU construction, the *CPFC* ensemble [52, 50], which is the product distribution over the following ensembles of n-qubit unitaries: P, a uniformly random permutation p on computational basis states, mapping $|x\rangle \mapsto |p(x)\rangle$; F, a random Boolean function implemented as a phase oracle $|x\rangle \mapsto (-1)^{f(x)}|x\rangle$; and C, the uniform distribution over n-qubit Clifford operators[6]. The C on the left and right are independent copies. These ensembles are statistically close to the Haar distribution, but can be efficiently instantiated as either t-designs or PRUs by replacing the P and F with pseudorandom components.

▶ **Theorem 2.11** ([50]). *The CPFC ensemble is a* ***strong*** $\varepsilon(n)$*-approximate* t*-design, with* $\varepsilon(n) = O(t^2/2^{n/8})$.

3 Random Unitaries in Constant (Quantum) Time

In this section, we analyze explicit constructions of the random unitary ensemble in Theorem 2.11. We show that each component of these constructions – the Clifford operator C, the phase oracle F, and the permutation P – can be instantiated by QAC^0_f circuits, which can then be bootstrapped into QAC^0 circuits. The latter step uses a recent technique for constructing random unitaries known as the *Gluing Lemma* [66]. To our knowledge, these are the first constructions of constant-depth unitary t-designs and PRUs with standard many-qubit gates [7].

3.1 Random Unitaries in QAC^0_f

First, we observe that both t-designs and pseudorandom unitaries can be implemented in QAC^0_f. To do so, it suffices to show that suitable derandomizations of P and F, as well as C, have QAC^0_f circuits.

3.1.1 Unitary Designs

To show that unitary designs can be implemented in QAC^0_f, we first consider implementing the Clifford operator C. A folklore construction from measurement-based quantum computing implements any Clifford circuit in constant depth, using mid-circuit measurements and classical feedforward. This construction can be easily translated into a QAC^0_f circuit [12].

[6] Or any unitary 2-design.

[7] As this work was finished, Zhang, Vijay, Gu, and Bao [78] gave a construction of unitary t-designs using depth using a Clifford operator followed by a depth $2^{O(t \log t)}$ circuit. In our constructions, we maintain constant depth, at the cost of increasing the number of ancillae.

▶ **Lemma 3.1.** *Let C be any n-qubit Clifford Circuit. C is implemented by a QAC^0_f circuit with $O(n^3)$ ancillae.*

Next, we consider the ensembles of phase oracles F. Just like random unitaries, a random boolean function requires exponentially many gates to implement, so suitable derandomizations must be chosen in order to obtain efficient circuits. For t-query adversaries, the following theorem from Zhandry [76] allows us to substitute a $2t$-wise independent function family for the oracle in F in place of a truly random function:

▶ **Lemma 3.2** ([76])**.** *Let Q be any quantum algorithm making at most t queries to F. The behavior of Q is unchanged if we replace F with a $2t$-wise independent function family (of phase oracles) $\overline{F}$.*

Moreover, $\overline{F}$ can be instantiated in TC^0 due to the following Theorem of Healy and Viola [36], who gave TC^0 circuits for a well-known family of t-wise independent functions:

▶ **Lemma 3.3** ([36])**.** *Let $\mathbb{F}_{2^s}$ be the finite field with 2^s elements. For any $t < 2^s$, the following is a family of t-wise independent functions:*

$$\overline{F} = \left\{ f(x) = \left(\sum_{i=0}^{t-1} a_i x^i \right) : (a_0, \ldots, a_{t-1}) \in \mathbb{F}_{2^s}^t \right\}$$

Furthermore, any function $f : \{0,1\}^s \to \{0,1\}$ from this family can be implemented by a TC^0 circuit with $\mathrm{poly}(s,t)$ gates.

Finally, we consider the random permutation P. To our knowledge, there are no known constructions of (even approximate) t-wise independent permutations in TC^0 for $t \in \omega(1)$ (for $t = O(1)$, the 2-round substitution-permutation network of Liu et al. [48] suffices). Fortunately, we can use the Luby-Rackoff inspired construction of Carolan and Cui et al. [13, 20]. In particular, [13] proved that the seven-round Feistel construction $\overline{P} = S_L S_R S_L S_R S_L S_R S_L$ is statistically indistinguishable from random up to inverse exponential error, even when inverse queries are allowed S_L and S_R are defined as follows:

$$S_L \left| x_1 \,\|\, x_2 \right\rangle = \left| x_1 \oplus f_L(x_2) \,\|\, x_2 \right\rangle, \quad S_R \left| x_1 \,\|\, x_2 \right\rangle = \left| x_1 \,\|\, x_2 \oplus f_R(x_1) \right\rangle$$

where f_L and f_R are random functions on $n/2$ bits. Again using Lemma 3.2, we can replace f_L and f_R with $2t$-wise independent functions using Lemma 3.3. Computing the required `XOR` of the first half of the input into the second can be implemented in $\mathsf{AC}^0 \subset \mathsf{TC}^0$, so $\overline{P}$ is implementable in TC^0 as desired. Putting together the previous claims in this section with Theorem 2.8, we conclude that any unitary from $C\overline{P}\overline{F}C$ can be implemented in QAC^0_f. Therefore,

▶ **Theorem 3.4.** *Strong ε-approximate unitary t-designs can be implemented by $\mathrm{poly}(n,t)$-size QAC^0_f circuits, with $\varepsilon = \exp(-\Omega(n))$.*

3.1.2 Pseudorandom Unitaries

To instantiate pseudorandom unitaries in QAC^0_f, we can repeat the analysis for unitary t-designs, using a quantum-secure pseudorandom function F^* used instead of a $2t$-wise independent function. A natural candidate F^* is the Ring-LWE pseudorandom function proposed by Banerjee, Peikert, and Rosen [5] and further analyzed by Zhandry [77]. This (candidate) quantum-secure pseudorandom function is keyed by a matrix $A \in \mathbb{Z}_p^{m \times n}$ and l different $n \times n$ matrices $\{S_i\}_{i=1}^l$ over $\mathbb{Z}_q$, and acts as follows:

$$\mathrm{PRF}_{(\mathbf{A},\{S_i\}_{i=1}^{l})}(x) = \left[\mathbf{A}^t \prod_{i=1}^{\ell} S_i^{x_i}\right]_p \tag{1}$$

For certain reasonable choices of p, q, m, l, [5] show that this PRF is secure and assuming that the Learning With Errors (LWE) [63] problem is post-quantum secure, a standard assumption in post-quantum cryptography. Moreover, [5] also show that Equation (1) can be typecast into a boolean function by using standard rounding techniques, while maintaining precision guarantees. Moreover, they also prove that this function can be implemented in TC^0:

▶ **Lemma 3.5** ([5, 77]). *Assuming post-quantum security of LWE The function in Equation* (1) *is $\varepsilon_0(n)$-secure against $t(n)$-time for quantum adversaries, the above function family is $\varepsilon_0(n)/\operatorname{poly}(n)$-secure against $t(n)$-time adversaries. Here, $\varepsilon_0(n)$ is a negligible function which depends the strength of the LWE assumption, and $t(n)$ is superpolynomial in n.*

Let P^* be defined analogously to $\overline{P}$ except with F^* in place of $\overline{F}$. Invoking the triangle inequality, the distinguishing advantage of any adversary that queries CP^*F^*C in place of a Haar-random unitary is at most the sum of the distinguishing advantage of F^*, P^*, and the error from Theorem 2.11:

▶ **Lemma 3.6.** *Suppose that F^* is $\varepsilon_0(n)$ secure against $t(n)$-time adversaries. Then, the ensemble CP^*F^*C is strongly $O(\varepsilon_0(n)) + \varepsilon(n))$ secure against $t(n)$-time adversaries, for $\varepsilon(n) = \operatorname{poly}(t(n))/\exp(-\Omega(n))$.*

Putting together the above claims again with Theorem 2.8, we conclude the following:

▶ **Theorem 3.7.** *Suppose the Learning With Errors (LWE) problem is $\varepsilon_0(n)$-secure against $t(n)$-time for quantum adversaries. Then, there exists an ensemble of unitaries implementable by $\operatorname{poly}(n)$-size QAC^0_f circuits, which is strongly $\Theta(\varepsilon_0(n)/\operatorname{poly}(n) + \operatorname{poly}(t(n))/\exp(-\Omega(n)))$-secure against $t(n)$-time adversaries.*

Taking t to be both subexponential and superpolynomial in n, and ε_0 to be any negligible function of n, we obtain an implementation of strong PRUs by QAC^0_f circuits. Moreover, assuming subexponential post-quantum security of LWE, i.e. that there exists some $c > 0$ such that $\varepsilon_0(n) = 1/2^{n^c}$-security holds all $t(n) = 2^{n^c}$-time algorithms, we obtain subexponentially secure strong PRUs by plugging in the relevant parameters.

▶ **Corollary 3.8.** *Suppose that LWE has subexponential post-quantum security*[8]*. Then, there exists an ensemble of unitaries implementable by $\operatorname{poly}(n)$-size QAC^0_f circuits, which is strongly $\Theta(1/2^{n^{c'}})$-secure against $2^{n^{c'}}$-time adversaries, for some $c' \in (0, 1)$.*

We use c' to emphasize that this constant may differ from the constant c in the preceding paragraph.

3.2 Random Unitaries in QAC^0

To bootstrap constructions of (pseudo)random unitaries in QAC^0_f to QAC^0, we will leverage the recently discovered "Gluing Lemma", proved by Schuster, Haferkamp, and Huang [66]. The Gluing Lemma is an extremely powerful statement, which allows random unitaries on n qubits to be "glued" from several copies of smaller $O(\operatorname{poly}\log(n))$ sized unitaries[9]. In particular, we have the two following statements, taken from and proved in [66]:

[8] This was the assumption in [66].

[9] A similar "gluing" result was proved by LaRacuente and Leditzky [45] to obtain low-depth unitary t-designs. For our purposes, it will be more convenient to use the constructions in [66]

▶ **Theorem 3.9** (Gluing Unitary t-designs, [66] Appendix B.3.)**.** *Consider a two layer, n-qubit brickwork circuit with local patch size ℓ. Suppose each gate in the brickwork is independently drawn from an ε/n-approximate unitary t-design. Then, as long as $\ell \geq \Theta(\log(nt^2/\varepsilon))$, the overall ensemble forms an ε-approximate unitary t-design on n qubits.*

▶ **Theorem 3.10** (Gluing Pseudorandom Unitaries, [66] Appendix C.4.)**.** *Consider a two layer, n-qubit brickwork circuit with local patch size ℓ. Suppose each ℓ-qubit gate in the brickwork is independently drawn from an ensemble which is $\varepsilon(\ell)$-secure against $t(\ell)$-time adversaries. Then, as long as $\ell \geq \omega(\log n)$, the overall ensemble is $O(n) \cdot \varepsilon(\ell)$-secure against $t(\ell)$-time adversaries.*

Since the brickwork circuit is only two layers, if the individual unitaries are implementable by QAC^0_f circuits, then the whole circuit is as well. Combining these with Theorem 3.4 and Theorem 3.7, we can exponentially reduce the width of the many-qubit gates required to form random unitaries (in our QAC^0_f constructions, the `FANOUT` gates act on $\mathrm{poly}(n)$ qubits):

▶ **Lemma 3.11.** *ε-approximate unitary t-designs can be implemented in QAC^0_f, where every `FANOUT` or `TOFFOLI` gate acts on at most $\mathrm{poly}(\log(nt^2/\varepsilon)) \cdot \mathrm{poly}(t)$ qubits.*

We have an analogous theorem for PRUs, combining Corollary 3.8 and Theorem 3.10:

▶ **Lemma 3.12.** *Let k be any constant. Suppose that LWE has subexponential post-quantum security. Taking a local patch size of $\ell = \log^{k/c} n$ in the statement of Theorem 3.10, we obtain an ensemble of unitaries which is $O(n) \cdot 1/2^{\log^{1/c} n}$-secure against $2^{\log^k n}$-time adversaries. Furthermore, every `FANOUT` or `TOFFOLI` gate acts on at most $\mathrm{poly} \log^k n$ qubits.*

Now that the width of the many-qubit gates has been reduced, our general strategy will be to substitute every `FANOUT` gate with an approximating QAC^0 circuit.

Rosenthal proved that `PARITY` (and therefore also `FANOUT`) is equivalent (up to a QAC^0 reduction to preparing a *n-nekomata* state, or a state of the form $\frac{|0^n,\psi_0\rangle+|1^n,\psi_1\rangle}{\sqrt{2}}$. A state is called an *$\varepsilon$-approximate n-nekomata* if it has fidelity at least $1-\varepsilon$ with some n-nekomata. In particular, Rosenthal showed the following:

▶ **Lemma 3.13.** *Let C be a circuit which constructs an ε-approximate nekomata from the all-zero state. Then, there exists a QAC^0 circuit, which makes queries to C and $C^\dagger$, which $1 - O(\varepsilon)$-implements the `PARITY` unitary operation.*

Rosenthal also showed that exponential-sized QAC^0 circuits exist, giving the first approximate implementation of `PARITY` in QAC^0 with exponentially many ancillae. Rosenthal's original construction was further analyzed by Grier and Morris [30], whose result we will use in what follows.

▶ **Fact 3.14** ([30])**.** *A w-qubit nekomata gate can be μ-approximated, with $\mu = 1/2^{\Theta(w)}$, by a constant-depth QAC^0 circuit C that uses $\exp(\Theta(w))$ many ancillae. Therefore, a width-w `PARITY` unitary operator can be $1/2^{\Theta(w)}$-implemented by a QAC^0 circuit with $\exp(\Theta(w))$ ancillae.*

This construction can be slightly improved in two ways while remaining in QAC^0. First, we observe that constant-sample amplification can be performed in QAC^0:

▶ **Corollary 3.15.** *For any constant d, a width-w `PARITY` unitary operator can be $1/2^{\Theta(dw)}$-implemented by a $O(\log d)$-depth QAC^0 circuit with $d \cdot \exp(\Theta(w))$ ancillae.*

Proof. Since QAC^0 contains reversible AND and OR gates with constant fan-in, d-qubit FANOUT and MAJORITY gates can be exactly implemented by QAC^0 circuits with $O(\log(d))$ ancillae. Therefore, as a preprocessing step to the Morris-Grier constructions, we first apply a width-$O(d)$ FANOUT gate to the input state. On each copy, we apply the Grier-Morris construction, take a majority vote of the PARITY computations, and uncompute. By a standard Chernoff bound, this decreases the error in the PARITY computation from $1/2^{\Theta(w)}$ to $1/2^{\Theta(dw)}$. ◀

Second, we use the downward self-reducibility of PARITY:

▶ **Corollary 3.16.** *For any constant d, a width-w^d PARITY unitary operator can be $w^d/2^{\Theta(dw)}$-implemented by a $O(d \log d)$-depth QAC^0 circuit with $dw^d \cdot \exp(\Theta(w))$ ancillae.*

Proof. PARITY is downward self-reducible: a width-w^d PARITY gate can be implemented by a depth-d tree where each leaf node corresponds to w^d qubits, and internal nodes which compute the PARITY of their w children. This tree has $O(w^d)$ nodes. If each internal node implements PARITY with error μ, then by Lemma 2.3 and Fact 3.14 the total error is at most $O(w^d\mu)$.

Since all internal nodes can be uncomputed after the root node computes the overall parity, the total depth increases by a multiplicative factor of $2d$, and each of the nodes uses at most $d \cdot \exp(O(w))$ ancillae by Corollary 3.15. Thus, we obtain a smaller circuit for implementing PARITY at the cost of increasing the depth and error parameters. ◀

Since PARITY and FANOUT are Hadamard conjugates (Fact 2.6), Corollary 3.16 also holds for FANOUT.

3.2.1 Unitary Designs

To obtain unitary t-designs in QAC^0, we replace each FANOUT gate in the "glued unitary" from Lemma 3.11 by an appropriate μ-implementation:

▶ **Theorem 3.17.** *For any $\varepsilon > 0$ and $c_0 \in \mathbb{N}$, ε-approximate n-qubit unitary t-designs can be μ-approximately implemented by size s, depth $O(c_1 \log c_1)$ QAC^0 circuits, with*

$$\mu = 1/\exp\left(\Theta\left(\frac{c_1 t}{c_0}\log(nt^2/\varepsilon)\right)\right) \quad s = \exp\left(\frac{t}{c_0}\log(nt^2/\varepsilon)\right)$$

Here, c_1 is any sufficiently large constant.

Proof. Lemma 3.11 allows us to construct an ε-approximate t-design using by implementing patches of smaller QAC^0_f unitaries each acting on $\ell = \mathrm{poly}(\log(nt^2/\varepsilon), t)$ qubits. Therefore, each FANOUT gate acts on at most ℓ qubits as well.

Let $c_0 \in \mathbb{N}$, and $w = \frac{t}{c_0}\log(nt^2/\varepsilon)$. Let $c_1 \in \mathbb{N}$ be such that $w^{c_1} \geq \ell$ – such a c_1 exists by Lemma 3.11). Corollary 3.16 implies that any FANOUT gate with width at most ℓ can be

$$w^{c_1}/2^{\Theta(c_1 w)} = 1/\exp\left(\Theta\left(\frac{c_1 t}{c_0}\log(nt^2/\varepsilon)\right)\right)$$

implemented by a QAC^0 circuit with

$$c_1 w^{c_1} \cdot \exp(w) = \exp\left(\frac{t}{c_0}\log(nt^2/\varepsilon)\right) \tag{2}$$

ancillae. Note that we have absorbed lower-order terms into $\Theta(\cdot)$ and $\exp(\cdot)$.

Since the original QAC^0_f circuit from Lemma 3.11 has at most $n\ell$ FANOUT gates, Lemma 2.3 implies that the overall implementation error is at most

$$\frac{n\ell}{\exp\left(\Theta\left(\frac{c_1 t}{c_0}\log(nt^2/\varepsilon)\right)\right)} = \frac{1}{\exp\left(\Theta\left(\frac{c_1 t}{c_0}\log(nt^2/\varepsilon)\right)\right)} \tag{3}$$

where we again absorbed lower-order terms into $\exp(\cdot)$. Similarly, the total number of ancillae required is at most $n\ell$ times the expression in Equation (2). ◀

A particularly important parameter regime is when t is polynomial in n, and ε is inverse polynomial in n. In this regime, we can μ-implement an ε-approximate t-design where μ is any inverse polynomial in n (by choosing c sufficiently large in Equation (3)). This construction requires $n \operatorname{poly}\log n \cdot n^{\delta t}$ ancillae, for any constant $\delta > 0$ (by choosing c_0 sufficiently large in Equation (2)). Hence, we obtain the following corollary:

▶ **Corollary 3.18.** *For any $\delta > 0$, and sufficiently large n depending on δ, there exists a QAC^0 circuit with $o(n^{1+\delta})$ ancillas which μ-implements an ε-approximate $O(1)$-design, where μ and ε are inverse polynomial in n.*

Note that the circuit depth scales as $O(1/\delta)$. Given that PARITY cannot be computed by depth-d QAC^0 circuits with $n^{1+1/3^d}$ ancillae [1], Corollary 3.18 implies that with slightly more ancillae ($n^{1+1/O(d)}$), we already obtain good approximations of constant t-designs (as long as n is sufficiently large). This observation will also hold for the pseudorandom unitary constructions in the next section.

Moreover, ε and μ can be both be made inverse-*quasipolynomial* in n, although the resulting constructions will require quasipolynomial ancillae. More generally, obtaining any construction with $o(1/\varepsilon)$ size would likely require improving the tradeoff between size and implementation error in either the Rosenthal or Morris-Grier constructions, which we leave as an interesting open problem.

3.2.2 Pseudorandom Unitaries

We can now implement PRUs in QAC^0 in an analogous fashion.

▶ **Theorem 3.19.** *For any $k, c_0 \in \mathbb{N}$, assuming subexponential post-quantum security of LWE, PRUs secure against $n^{\operatorname{poly}\log n}$-time adversaries can be μ-implemented by size s, depth $O(c\log c)$ QAC^0 circuits, with*

$$\mu = 1/\exp(\Theta(c\log^k n/c_0)) \quad s = \exp(\log^k n/c_0)$$

Again, c is any sufficiently large constant (where "sufficiently large" depends on c_0).

Proof. We reuse the analysis of Theorem 3.19, with $w = \log^k n/c_0$ and local patch size $\ell = w^c$ for some constant c (the ℓ from Lemma 3.11), obtaining a $1/\exp(\Theta(c\log n/c_0))$-implementation with $\exp(\log n/c_0)$ ancillae. ◀

Taking $k = 1$, we obtain implementations of PRUs by polynomial-size QAC^0 circuits. Like Corollary 3.18, the polynomial can be chosen to be $o(n^{1+\delta})$ for any $\delta > 0$, almost linear when $\log n \gg c_0$. When $k > 1$, we implement PRUs to inverse *quasipolynomial* error, although the resulting circuits have quasipolynomial size.

One important caveat implicit in both QAC^0 constructions is that the desired random unitary on the n input qubits is only approximately implemented, and there is no guarantee that the action when tracing out the ancilla qubits is unitary[10]. In many applications, this scenario is addressed by giving the distinguishing algorithm access to the *channel* $\Phi_C(\rho) := \mathrm{Tr}_{anc}(C(\rho \otimes |0^a\rangle\langle 0^a|)C^\dagger)$, which maps the state of the n input qubits to the reduced state on the same n qubits after applying C.

If C μ-implements a unitary U (on n qubits), then replacing queries to U with queries to Φ_C incurs an additive error:

▶ **Lemma 3.20.** *Assume C ε-approximates a unitary U. Let Q be any quantum algorithm which makes at most t queries to either $\Phi_C/\Phi_{C^\dagger}$ or $U/U^\dagger$. For any input state ρ,*

$$\left\| Q^{\Phi_C,\Phi_{C^\dagger}}\rho(Q^{\Phi_C,\Phi_{C^\dagger}})^\dagger - Q^{U,U^\dagger}\rho(Q^{U,U^\dagger})^\dagger \right\|_{td} \leq t\mu$$

Proof. We assume without loss of generality that all queries are to U: the proof is identical if queries to $U^\dagger$ are replaced by queries to $\Phi_{C^\dagger}$. Define $Q_0 = Q^{\Phi_C}, Q_t = Q^U$, and Q_i the algorithm whose first i queries are to U and whose last $t-i$ queries are to Φ_C. By the triangle inequality, this is at most

$$\begin{aligned} & t \cdot \sum_{i=0}^{t-1} \left\| Q_{i+1}\rho(Q_{i+1})^\dagger - Q_i\rho(Q_i)^\dagger \right\|_{td} \\ \leq\ & t \cdot \sum_{i=0}^{t-1} \left\| \mathcal{E}_i(U\rho_i U^\dagger) - \mathcal{E}_i(\mathrm{Tr}_{anc}(C(\rho_i \otimes |0^a\rangle\langle 0^a|)C^\dagger)) \right\|_{td} \\ \leq\ & t \cdot \sum_{i=0}^{t-1} \left\| U\rho_i U^\dagger - \mathrm{Tr}_{anc}(C(\rho_i \otimes |0^a\rangle\langle 0^a|)C^\dagger) \right\|_{td} \end{aligned}$$

for some channels $\mathcal{E}_i$ and states ρ_i, where in the second line we use that trace distance can only decrease under an application of $\mathcal{E}_i$. Writing ρ_i as a linear combination $\rho_i = \sum_j \alpha_{ij} |\psi_j\rangle\langle\psi_j|$,

$$\begin{aligned} & t \cdot \sum_{i=0}^{t-1} \sum_j \alpha_{ij} \left\| U|\psi_j\rangle\langle\psi_j|U^\dagger - \mathrm{Tr}_{anc}(C(|\psi_j\rangle\langle\psi_j| \otimes |0^a\rangle\langle 0^a|)C^\dagger) \right\|_{td} \\ \leq\ & t \cdot \sum_{i=0}^{t-1} \sum_j \alpha_{ij} \left\| U|\psi_j\rangle\langle\psi_j|U^\dagger \otimes |0^a\rangle\langle 0^a| - C(|\psi_j\rangle\langle\psi_j| \otimes |0^a\rangle\langle 0^a|)C^\dagger \right\|_{td} \\ \leq\ & t \cdot \sum_{i=0}^{t-1} \sum_j \alpha_{ij}(1 - \sqrt{1-\mu}) = t \cdot (1 - \sqrt{1-\mu}) \leq t\mu \end{aligned}$$

from the first to the second line, we used that trace distance is non-increasing when taking a partial trace. From the second to third line, we use the Fuchs-van-de Graff inequality and the definition of μ-implementation. ◀

Since the maximum distinguishing advantage between two quantum states is equal to the trace distance, combining Theorem 3.19 and Lemma 3.20 implies that any $n^{\mathrm{poly}\log n}$-time algorithm which makes at most t queries to a polynomial size QAC^0 circuit C cannot, in the worst case distinguish C from Haar-random except with at most t times that probability. In particular, we obtain the following distinguishing lower bound:

[10] In fact, it is an open problem to exactly implement `FANOUT` using QAC^0 circuits of any size.

▶ **Theorem 3.21.** *Assuming subexponential post-quantum security of LWE, no fixed polynomial time algorithm can distinguish the following two cases with non-negligible probability:*

- *C is a polynomial size QAC^0 circuit.*
- *$\Phi_C = \Phi_U = U\rho U^\dagger$ for a Haar-random unitary U.*

Proof. Assume by contradiction that a $O(n^{c_0})$-time algorithm A could distinguish any polynomial-size QAC^0 circuit from random, with probability $\Omega(1/n^{c_1})$. Use Theorem 3.19 to implement a PRU to error $O(1/n^{c_2})$, where $c_2 \gg c_0 + c_1$ in polynomial-size QAC^0. Then, Lemma 3.20 implies that A's maximum distinguishing advantage should be $O(n^{c_0-c_2}) = o(1/n^{c_1})$, a contradiction. ◀

3.3 Learning Lower-Bounds from Pseudorandomness

We will now show how our implementations of PRUs (and quantum secure PRFs as in Equation (1)) can be used to prove powerful computational lower bounds for average-case learning of QAC^0.

In particular, we will prove bounds with respect to the average-case distance measure, which is defined for two n-qubit unitary channels as

$$\mathcal{D}_{\text{avg}}(U,V) = \mathop{\mathbb{E}}_{|\psi\rangle\sim\mu_{2^n}} \left[1 - F\left(U(|\psi\rangle\langle\psi|U^\dagger), V(|\psi\rangle\langle\psi|V^\dagger)\right)\right],$$

where $|\psi\rangle$ is sampled from the Haar measure μ_{2^n} and $F(\cdot,\cdot)$ is the fidelity. We begin by using the PRU construction of Theorem 3.19 (assuming hardness of Learning with Errors) to prove a super-polynomial computational lower bound for learning QAC^0 circuits with polynomial ancillae.

▶ **Lemma 3.22.** *Assuming subexponential post-quantum security of LWE, $n^{\omega(1)}$ time is necessary to learn polynomial-sized QAC^0 unitaries, according to the average-case distance measure.*

Proof. By contradiction, assume there exists an efficient algorithm for learning a circuit $\widetilde{U}$. We show that there exists a polynomial time algorithm for distinguishing U from a Haar random unitary, via the learned circuit $\widetilde{U}$, contradicting the definition of a PRU. For ease of notation, we will denote $d = 2^n$. First, we convert each unitary into a state via the Choi-Jamiolkowski isomorphism. In particular, let $|\Phi\rangle = \frac{1}{\sqrt{d}}\sum_{i=1}^{d} |i\rangle \otimes |i\rangle$ denote the maximally entangled state, such that $|\psi_U\rangle = (U \otimes I)|\Phi\rangle$ and $|\psi_{\widetilde{U}}\rangle = (\widetilde{U} \otimes I)|\Phi\rangle$ are the respective Choi states of U and $\widetilde{U}$. Performing a swap test between $|\psi_U\rangle$ and $|\psi_{\widetilde{U}}\rangle$ results in the zero outcome with probability

$$\begin{aligned} P_0 &= \frac{1}{2}\left(1 + |\langle\psi_{\widetilde{U}}|\psi_U\rangle|^2\right) = \frac{1}{2}\left(1 + \left|\frac{1}{d}\sum_{i,j}\langle j|\widetilde{U}^\dagger U|i\rangle\langle j|i\rangle\right|^2\right) \\ &= \frac{1}{2}\left(1 + \frac{1}{d^2}\left|\sum_i \langle i|\widetilde{U}^\dagger U|i\rangle\right|^2\right) = \frac{1}{2}\left(1 + \frac{1}{d^2}|\text{Tr}(\widetilde{U}^\dagger U)|^2\right). \end{aligned}$$

In the case that unitary U is a QAC^0 PRU, by the guarantees of the learning algorithm, this implies that

$$\mathcal{D}_{\text{avg}}(U,\widetilde{U}) = \frac{d}{d+1}\left(1 - \frac{1}{d^2}|\text{Tr}(\widetilde{U}^\dagger U)|^2\right) \le \varepsilon \implies P_0 \ge 1 - \varepsilon\cdot\frac{d+1}{2d},$$

for $\varepsilon = 1/\operatorname{poly}(n)$, meaning that $P_0 \approx 1$. Meanwhile, in the case that U is Haar-random, for any fixed (learned) $\widetilde{U}$, the expectation of P_0 is tightly concentrated around

$$\mathbb{E}_{U\sim\mu}[P_0] = \frac{1}{2}\left(1 + \frac{1}{d^2}\mathbb{E}_{U\sim\mu}\left[|\mathrm{Tr}(\widetilde{U}^\dagger U)|^2\right]\right) = \frac{1}{2} + \frac{1}{2d^2} \approx \frac{1}{2}.$$

Therefore, by running the swap test just a constant number of times, we can estimate $\hat{P}_0$ and decide, say, if $\hat{P}_0 < 3/4$ then U is Haar random (otherwise it is not) – thus distinguishing PRUs from Haar random unitaries. ◀

Using an analogous proof approach and just the PRFs described in Section 3.1.2 (through which the previous PRUs were constructed), we are also able to extend the lower-bound to the sub-linear ancilla setting (improved over the polynomial ancilla required for the PRU implementation). This proof structure is more directly analogous to the proofs of classical [42] and quantum [3] hardness for average-case learning of AC^0.

▶ **Corollary 3.23.** *Fix any constant $\delta > 0$. Assuming subexponential post-quantum security of LWE, $n^{\omega(1)}$ time is necessary to learn QAC^0 unitaries with $O(n^\delta)$ ancillae, according to the average-case distance measure.*

Note that the ancilla-overhead reduction achieved by Corollary 3.23 relative to Lemma 3.22 is not fundamental, but rather a reflection of PRU implementation's reliance on PRFs, with additional ancilla overhead. However, we still believe the PRU proof given in Lemma 3.22 is of interest, as there could plausibly exist ancilla-efficient PRU constructions that are not based on PRFs.

▶ **Corollary 3.24.** *Assuming the quantum LMN low-support conjecture [73, Conjecture 1], the quasipolynomial $O(n^{\operatorname{poly}\log n})$ time-complexity of the [73] algorithm nearly matches our super-polynomial $n^{\omega(1)}$ time lower-bound and is, thus, near-optimal.*

Although an exponential worst-case learning lower-bound was previously known for QAC^0 [39, 73], this is the first *average-case* lower-bound.

3.4 Measurement-Based Random Unitaries

As a consequence of the equivalence between Model 2 and Model 3 from Section 1, constructions of random unitaries in QAC^0_f can be recast as constant-time *measurement-based* implementations of random unitaries:

▶ **Theorem 3.25.** *Using constant-depth quantum circuits with two-dimensional nearest-neighbor gates and feedforward measurement, we can exactly implement*

- *Standard ε-approximate t-designs with $n \cdot \operatorname{poly}(\log nt^2/\varepsilon) \cdot \operatorname{poly}(t)$ ancillae, and quasipolynomial-secure PRUs (assuming subexponential security of LWE) with $n \operatorname{poly}\log n$ ancillae.*
- *Strong t-designs (to inverse exponential error) with $\operatorname{poly}(n,t)$ ancillae, and strong PRUs with $\operatorname{poly}(n)$ ancillae.*

Proving 3.25 amounts to applying measurement-based techniques to parallelize layers of `FANOUT` gates in 3.4, 3.7. Geometric locality can be enforced using `SWAP` gates, which are also Clifford and can be parallelized using 3.1.

3.5 Towards $\texttt{PARITY} \notin \mathsf{QAC}^0$

In this section, we show how the construction of strong unitary t-designs in QAC^0_f implies a connection between random unitaries and $\texttt{PARITY} \notin \mathsf{QAC}^0$. Specifically, we show that any algorithm which distinguishes QAC^0 circuits from Haar-random unitaries using a polynomial number of forward- and inverse-queries implies that $\texttt{PARITY} \notin \mathsf{QAC}^0$:

▶ **Theorem 3.26.** *Suppose that there exists a quantum algorithm which makes* $\mathrm{poly}(n)$ *queries to* Φ_C *and* $\Phi_{C^\dagger}$*, and distinguishes the following two cases with at least inverse exponential advantage:*

- C *is a* QAC^0 *circuit.*
- $\Phi_C = \Phi_U = U\rho U^\dagger$ *for a Haar-random unitary* U*.*

Then, no QAC^0 *circuit* C' *computes the boolean function* $\texttt{PARITY}(x) = \bigoplus_i x_i$ *exactly.*

Theorem 3.26 may unlock alternate paths towards proving $\texttt{PARITY} \notin \mathsf{QAC}^0$. For illustration, we show that having a single "peak" in the Pauli decomposition of a Heisenberg-evolved observable suffices to distinguish a unitary from Haar-random. In general, this requires inverse access to U. Variants of this property are implied by [54, 1] – we leave details for future exploration. To illustrate, we give an example of a candidate property which, if true, would separate QAC^0 and QAC^0_f. Suppose that for any unitary U implemented by a QAC^0 circuit C, we have

$$\left|\mathrm{Tr}(UOU^\dagger P)/2^n\right|^2 \gg 1/\exp(n)$$

for some Pauli string P, with high probability over a single qubit operator O[11]. Using standard Pauli sampling techniques, we can sample a Pauli string P' with probability $\left|\mathrm{Tr}(UOU^\dagger P')/2^n\right|^2$. For a Haar-random U, $\left|\mathrm{Tr}(UOU^\dagger P)/2^n\right|^2$ will almost certainly be exponentially small for all P. Thus, this test distinguishes U from a Haar-random unitary, without appealing to a notion of low-degree concentration.

Proof of Theorem 3.26. Suppose for the sake of contradiction that a QAC^0 circuit U exactly computes the (boolean function) $\texttt{PARITY}$, so U is an $(n+a)$-qubit unitary taking n input qubits and $a = \mathrm{poly}(n)$ ancillae such that

$$U\left|x_1,\ldots,x_n,0\right\rangle\left|0^a\right\rangle = \left|\texttt{PARITY}(x)\right\rangle\left|\psi_x\right\rangle.$$

For some $n+a-1$-qubit state $|\psi_x\rangle$. Note that if we trace out the ancilla qubits, this channel is not in general unitary. We would like to use this to construct the *unitary* implementation of $\texttt{PARITY}$. This can be done using the standard uncomputation technique: on input $|x\rangle$ we apply U to $|x\rangle|0^a\rangle$ and CNOT the value $\texttt{PARITY}(x)$ into a separate register in the state $|b\rangle$. We then uncompute by applying $U^\dagger$ to get $|x, \texttt{PARITY}(x)\oplus b\rangle|0^a\rangle$. Therefore we can construct a circuit U' that implements the $\texttt{PARITY}$ unitary $U'|x_1,\ldots,x_n,b\rangle|0^a\rangle = \texttt{PARITY}|x_1,\ldots,x_n,b\rangle\otimes|0^a\rangle$. Since we only used $U, U^\dagger$ and a CNOT gate, U' is also a QAC^0 circuit.

Using the ability to compute $\texttt{PARITY}$ together with the construction of strong unitary t-designs in Theorem 3.4, we conclude that for any $t = \mathrm{poly}(n)$, strong $\exp(-\Omega(n))$-approximate unitary t-designs can be implemented in QAC^0. Together with 3.20, this contradicts the theorem assumptions. ◀

[11] This is a weaker statement than the "low-degree concentration" conjectures of [54, 73].

Furthermore, if there exists an algorithm which distinguishes QAC^0 circuits from Haar-random unitaries with a polynomial number of queries, then this implies that QAC^0 circuits cannot compute the *classical parity function*, even in the average-case. To do so, one can show that the ability to compute a classical parity function implies that the standard measurement-based construction of a *nekomata state* [64], which allows one to synthesize the PARITY unitary. Using 3.4, this would give strong $\exp(-\Omega(n))$-approximate unitary $\mathrm{poly}(n)$-designs can be approximated to arbitrary inverse polynomial precision in QAC^0, a contradiction.

Together with theorem Theorem 3.26, we have demonstrated connect random unitaries and quantum circuit lower bounds. This connection was previously studied in the context of quantum learning of boolean functions [2], and for state synthesis [16]. The fact that distinguishing (channels approximating) QAC^0 *unitaries* suffices to prove quantum circuit lower bounds for *decision problems* seems to be a particular feature of QAC^0, enabled by the QAC^0 reduction of implementing PARITY to synthesizing nekomata states (Lemma 3.13).

References

1 Anurag Anshu, Yangjing Dong, Fengning Ou, and Penghui Yao. On the Computational Power of QAC^0 with Barely Superlinear Ancillae, 2024. doi:10.48550/arXiv.2410.06499.

2 Srinivasan Arunachalam, Alex B. Grilo, Tom Gur, Igor C. Oliveira, and Aarthi Sundaram. Quantum learning algorithms imply circuit lower bounds, 2021. doi:10.48550/arXiv.2012.01920.

3 Srinivasan Arunachalam, Alex Bredariol Grilo, and Aarthi Sundaram. Quantum hardness of learning shallow classical circuits. *SIAM Journal on Computing*, 50(3):972–1013, 2021. doi:10.1137/20M1344202.

4 Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando Brandao, David Buell, Brian Burkett, Yu Chen, Jimmy Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Michael Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew Harrigan, Michael Hartmann, Alan Ho, Markus Rudolf Hoffmann, Trent Huang, Travis Humble, Sergei Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, Dave Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod Ryan McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin Jeffery Sung, Matt Trevithick, Amit Vainsencher, Benjamin Villalonga, Ted White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574:505–510, 2019. URL: https://www.nature.com/articles/s41586-019-1666-5.

5 Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. Cryptology ePrint Archive, Paper 2011/401, 2011. URL: https://eprint.iacr.org/2011/401.

6 Jinge Bao and Francisco Escudero-Gutiérrez. Learning junta distributions, quantum junta states, and QAC^0 circuits, 2025. doi:10.48550/arXiv.2410.15822.

7 Ron Belyansky, Przemyslaw Bienias, Yaroslav A Kharkov, Alexey V Gorshkov, and Brian Swingle. Minimal model for fast scrambling. *Physical review letters*, 125(13):130601, 2020.

8 Avrim Blum, Merrick Furst, Michael Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In *Advances in Cryptology – CRYPTO 1994*, volume 839 of *Lecture Notes in Computer Science*, pages 278–291. Springer, 1994. doi:10.1007/3-540-48329-2_24.

9 Dolev Bluvstein, Harry Levine, Giulia Semeghini, Tout T. Wang, Sepehr Ebadi, Marcin Kalinowski, Alexander Keesling, Nishad Maskara, Hannes Pichler, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604(7906):451–456, April 2022. doi:10.1038/s41586-022-04592-6.

10 Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346(2):397–434, August 2016. doi:10.1007/s00220-016-2706-8.

11 Mark Braverman. Polylogarithmic independence fools AC^0 circuits. *Journal of the ACM (JACM)*, 57(5):1–10, 2008.

12 Harry Buhrman, Marten Folkertsma, Bruno Loff, and Niels M. P. Neumann. State preparation by shallow circuits using feed forward. *Quantum*, 8:1552, December 2024. doi:10.22331/q-2024-12-09-1552.

13 Joseph Carolan. Compressed permutation oracles. Cryptology ePrint Archive, Paper 2025/1734, 2025. URL: https://eprint.iacr.org/2025/1734.

14 Kean Chen, Qisheng Wang, and Zhicheng Zhang. Local test for unitarily invariant properties of bipartite quantum states, 2025. doi:10.48550/arXiv.2404.04599.

15 Yanlin Chen, Yilei Chen, Rajendra Kumar, Subhasree Patro, and Florian Speelman. Fine-grained complexity via quantum natural proofs, 2025. doi:10.48550/arXiv.2504.10363.

16 Nai-Hui Chia, Daniel Liang, and Fang Song. Quantum state learning implies circuit lower bounds, 2024. doi:10.48550/arXiv.2405.10242.

17 Richard Cleve and John Watrous. Fast parallel circuits for the quantum fourier transform, 2000. arXiv:quant-ph/0006004.

18 Antonio D Córcoles, Maika Takita, Ken Inoue, Scott Lekuch, Zlatko K Minev, Jerry M Chow, and Jay M Gambetta. Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits. *Physical Review Letters*, 127(10):100501, 2021.

19 Jordan Cotler, Thomas Schuster, and Masoud Mohseni. Information-theoretic hardness of out-of-time-order correlators, 2022. doi:10.48550/arXiv.2208.02256.

20 Laura Cui, Thomas Schuster, Fernando Brandao, and Hsin-Yuan Huang. Unitary designs in nearly optimal depth, 2025. doi:10.48550/arXiv.2507.06216.

21 Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1), July 2009. doi:10.1103/physreva.80.012304.

22 Mark R. Dowling and Michael A. Nielsen. The geometry of quantum computation, 2006. arXiv:quant-ph/0701004.

23 Joseph Emerson, Yaakov S. Weinstein, Marcos Saraceno, Seth Lloyd, and David G. Cory. Pseudo-Random Unitary Operators for Quantum Information Processing. *Science*, 302(5653):2098–2101, December 2003. doi:10.1126/science.1090790.

24 Maosen Fang, Stephen Fenner, Frederic Green, Steven Homer, and Yong Zhang. Quantum lower bounds for fanout, 2003. arXiv:quant-ph/0312208.

25 Stephen Fenner, Daniel Grier, Daniel Padé, and Thomas Thierauf. Tight bounds on depth-2 QAC-circuits computing parity, 2025. doi:10.48550/arXiv.2504.06433.

26 Stephen A. Fenner. Implementing the fanout gate by a Hamiltonian, 2003. arXiv:quant-ph/0309163.

27 Austin G Fowler. Time-optimal quantum computation. *arXiv preprint arXiv:1210.4626*, 2012.

28 M. Furst, J. Saxe, and M. Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984. doi:10.1007/BF01744431.

29 Alaina M Green, A Elben, C Huerta Alderete, Lata Kh Joshi, Nhung H Nguyen, Torsten V Zache, Yingyue Zhu, Bhuvanesh Sundar, and Norbert M Linke. Experimental measurement of out-of-time-ordered correlators at finite temperature. *Physical Review Letters*, 128(14):140601, 2022.

30 Daniel Grier and Jackson Morris. Quantum threshold is powerful, 2024. doi:10.48550/arXiv.2411.04953.

31 Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1579–1588, 2023. doi:10.1145/3564246.3585198.

32 Tom Gur, Min-Hsiu Hsieh, and Sathyawageeswar Subramanian. Sublinear quantum algorithms for estimating von neumann entropy, 2021. doi:10.48550/arXiv.2111.11139.

33 Jonas Haferkamp. Random quantum circuits are approximate unitary t-designs in depth $o(nt^5 + o(1))$. *Quantum*, 6:795, September 2022. doi:10.22331/q-2022-09-08-795.

34 Aram W. Harrow and Saeed Mehraban. Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates. *Communications in Mathematical Physics*, 401(2):1531–1626, May 2023. doi:10.1007/s00220-023-04675-z.

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

36 Alexander Healy and Emanuele Viola. Constant-depth circuits for arithmetic in finite fields of characteristic two. In *Proceedings of the 23rd Annual Conference on Theoretical Aspects of Computer Science*, STACS'06, pages 672–683, Berlin, Heidelberg, 2006. Springer-Verlag. doi:10.1007/11672142_55.

37 Peter Hoyer and Robert Spalek. Quantum fan-out is powerful. *Theory of Computing*, 1(1):81–103, 2005. doi:10.4086/toc.2005.v001a005.

38 Hsin-Yuan Huang, Sitan Chen, and John Preskill. Learning to predict arbitrary quantum processes, 2023. doi:10.48550/arXiv.2210.14894.

39 Hsin-Yuan Huang, Yunchao Liu, Michael Broughton, Isaac Kim, Anurag Anshu, Zeph Landau, and Jarrod R. McClean. Learning Shallow Quantum Circuits. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, pages 1343–1351, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3618260.3649722.

40 Mohsin Iqbal, Nathanan Tantivasadakarn, Ruben Verresen, Sara L. Campbell, Joan M. Dreiling, Caroline Figgatt, John P. Gaebler, Jacob Johansen, Michael Mills, Steven A. Moses, Juan M. Pino, Anthony Ransford, Mary Rowe, Peter Siegfried, Russell P. Stutz, Michael Foss-Feig, Ashvin Vishwanath, and Henrik Dreyer. Non-abelian topological order and anyons on a trapped-ion processor. *Nature*, 626(7999):505–511, February 2024. doi:10.1038/s41586-023-06934-4.

41 Sven Jandura, Vineesha Srivastava, Laura Pecorari, Gavin K. Brennen, and Guido Pupillo. Nonlocal multiqubit quantum gates via a driven cavity. *Physical Review A*, 110(6), December 2024. doi:10.1103/physreva.110.062610.

42 Michael Kharitonov. Cryptographic hardness of distribution-specific learning. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '93, pages 372–381, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/167088.167197.

43 Isaac H. Kim. Catalytic z-rotations in constant t-depth, 2025. doi:10.48550/arXiv.2506.15147.

44 Robbie King, Kianna Wan, and Jarrod R. McClean. Exponential learning advantages with conjugate states and minimal quantum memory. *PRX Quantum*, 5(4), October 2024. doi:10.1103/prxquantum.5.040301.

45 Nicholas LaRacuente and Felix Leditzky. Approximate unitary k-designs from shallow, low-communication circuits, 2024. doi:10.48550/arXiv.2407.07876.

46 Zeyang Li, Boris Braverman, Simone Colombo, Chi Shu, Akio Kawasaki, Albert F Adiyatullin, Edwin Pedrozo-Peñafiel, Enrique Mendez, and Vladan Vuletić. Collective spin-light and light-mediated spin-spin interactions in an optical cavity. *PRX Quantum*, 3(2):020308, 2022.

47 Nathan Linial, Yishay Mansour, and Noam Nisan. Constant depth circuits, fourier transform, and learnability. *J. ACM*, 40(3):607–620, July 1993. doi:10.1145/174130.174138.

48 Tianren Liu, Angelos Pelecanos, Stefano Tessaro, and Vinod Vaikuntanathan. Layout graphs, random walks and the t-wise independence of SPN block ciphers. Cryptology ePrint Archive, Paper 2024/083, 2024. doi:10.1007/978-3-031-38548-3_23.

49 Chuhan Lu, Minglong Qin, Fang Song, Penghui Yao, and Mingnan Zhao. Parallel kac's walk generates pru, 2025. doi:10.48550/arXiv.2504.14957.

50 Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries, 2024. doi:10.48550/arXiv.2410.10116.

51 Easwar Magesan, J. M. Gambetta, and Joseph Emerson. Scalable and robust randomized benchmarking of quantum processes. *Physical Review Letters*, 106(18), May 2011. doi:10.1103/physrevlett.106.180504.

52 Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Simple constructions of linear-depth t-designs and pseudorandom unitaries, 2024. doi:10.48550/arXiv.2404.12647.

53 Cristopher Moore. Quantum circuits: Fanout, parity, and counting, 1999. arXiv:quant-ph/9903046.

54 Shivam Nadimpalli, Natalie Parham, Francisca Vasconcelos, and Henry Yuen. On the Pauli Spectrum of QAC^0. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC '24, pages 1498–1506. ACM, June 2024. doi:10.1145/3618260.3649662.

55 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

56 Emilio Federico Onorati. *Random Processes over the Unitary Group: Mixing Properties and Applications in Quantum Information*. Freie Universitaet Berlin (Germany), 2019.

57 Daniel Padé, Stephen Fenner, Daniel Grier, and Thomas Thierauf. Depth-2 QAC circuits cannot simulate quantum parity. *arXiv preprint arXiv:2005.12169*, 2020. arXiv:2005.12169.

58 Natalie Parham. Quantum circuit lower bounds in the magic hierarchy, 2025. doi:10.48550/arXiv.2504.19966.

59 Avikar Periwal, Eric S Cooper, Philipp Kunkel, Julian F Wienand, Emily J Davis, and Monika Schleier-Smith. Programmable interactions and emergent geometry in an array of atom clouds. *Nature*, 600(7890):630–635, 2021.

60 Ran Raz and Avishay Tal. Oracle separation of BQP and PH. *ACM Journal of the ACM (JACM)*, 69(4):1–21, 2022. doi:10.1145/3530258.

61 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*, 41:333–338, 1987. URL: https://api.semanticscholar.org/CorpusID:121744639.

62 Alexander A. Razborov and Steven Rudich. Natural proofs, 1997. 26th Annual ACM Symposium on the Theory of Computing (STOC '94) (Montreal, PQ, 1994). doi:10.1006/jcss.1997.1494.

63 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography, 2024. doi:10.48550/arXiv.2401.03703.

64 Gregory Rosenthal. Bounds on the QAC^0 Complexity of Approximating Parity. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.48550/arXiv.2008.07470.

65 Anish Saini, Athanasios Tsokanos, and Raimund Kirner. Quantum randomness in cryptography – a survey of cryptosystems, RNG-based ciphers, and QRNGs. *Information*, 13(8):358, 2022. doi:10.3390/info13080358.

66 Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. Random unitaries in extremely low depth, 2025. doi:10.48550/arXiv.2407.07754.

67 Peter W Shor. Scheme for reducing decoherence in quantum computer memory. *Physical review A*, 52(4):R2493, 1995.

68 Kevin C. Smith, Abid Khan, Bryan K. Clark, S.M. Girvin, and Tzu-Chieh Wei. Constant-depth preparation of matrix product states with adaptive quantum circuits. *PRX Quantum*, 5(3), September 2024. doi:10.1103/prxquantum.5.030344.

69 Yongxin Song, Liberto Beltrán, Ilya Besedin, Michael Kerschbaum, Marek Pechal, François Swiadek, Christoph Hellings, Dante Colao Zanuz, Alexander Flasby, Jean-Claude Besse, and Andreas Wallraff. Realization of constant-depth fan-out with real-time feedforward on a superconducting quantum processor, 2024. doi:10.48550/arXiv.2409.06989.

70 Florian Speelman. Instantaneous Non-Local Computation of Low T-Depth Quantum Circuits. In *11th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2016, Berlin, Germany, September 27-29, 2016*, volume 61 of *LIPIcs*, pages 9:1–9:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICS.TQC.2016.9.

71 Yasuhiro Takahashi and Seiichiro Tani. Collapse of the hierarchy of constant-depth exact quantum circuits, 2012. doi:10.48550/arXiv.1112.6063.

72 Nathanan Tantivasadakarn, Ashvin Vishwanath, and Ruben Verresen. Hierarchy of topological order from finite-depth unitaries, measurement, and feedforward. *PRX Quantum*, 4(2):020339, 2023.

73 Francisca Vasconcelos and Hsin-Yuan Huang. Learning shallow quantum circuits with many-qubit gates, 2025. doi:10.48550/arXiv.2410.16693.

74 Joel J. Wallman and Joseph Emerson. Noise tailoring for scalable quantum computation via randomized compiling. *Physical Review A*, 94(5), November 2016. doi:10.1103/physreva.94.052325.

75 Dongmin Yu, Han Wang, Jin ming Liu, Shi-Lei Su, Jing Qian, and Weiping Zhang. Multiqubit toffoli gates and optimal geometry with rydberg atoms, 2022. doi:10.48550/arXiv.2203.14302.

76 Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. *International Journal of Quantum Information*, 13(04):1550014, 2015. doi:10.1142/S0219749915500148.

77 Mark Zhandry. How to construct quantum random functions. *J. ACM*, 68(5), August 2021. doi:10.1145/3450745.

78 Yuzhen Zhang, Sagar Vijay, Yingfei Gu, and Yimu Bao. Designs from magic-augmented Clifford circuits, 2025. doi:10.48550/arXiv.2507.02828.

A Appendix

▶ **Lemma A.1.** *Assume that C ε_0-approximates U and D ε_1-approximates V. Then, CD $(\varepsilon_0+\varepsilon_1)$-approximates UV.*

Proof. We want to lower bound

$$\left|\langle\psi,0^a|\, D^\dagger C^\dagger (UV\,|\psi\rangle\otimes|0^a\rangle)\right|^2$$

Let $|\phi\rangle = V\,|\psi\rangle\otimes|0^a\rangle$. Using linearity of the inner product, we obtain the following chain of inequalities:

$$\begin{aligned}
\left|\langle\psi,0^a|\, D^\dagger C^\dagger (U\otimes I)\,|\phi\rangle\right| &= \left|\langle\psi,0^a|\, D^\dagger\,|\phi\rangle\right| + \left|\langle\psi,0^a|\, D^\dagger (C^\dagger(U\otimes I)-I)\,|\phi\rangle\right| \\
\geq \sqrt{(1-\varepsilon_1)} + \left|\langle\psi,0^a|\, D^\dagger (C^\dagger(U\otimes I)-I)\,|\phi\rangle\right| &\geq \sqrt{(1-\varepsilon_1)} + \left\|C^\dagger(U\otimes I)\,|\phi\rangle - |\phi\rangle\right\|_2 \\
&= \sqrt{(1-\varepsilon_1)} + \sqrt{1-\varepsilon_0} - 1 \\
\implies \left|\langle\psi,0^a|\, D^\dagger C^\dagger (UV\,|\psi\rangle\otimes|0^a\rangle)\right|^2 &\geq \left(\sqrt{(1-\varepsilon_1)} + \sqrt{1-\varepsilon_0} - 1\right)^2 \geq 1-\varepsilon_0-\varepsilon_1
\end{aligned}$$

from the first to the second line, we used that D ε_1-approximates V, and from the third to fourth line that C ε_0-approximates U. ◀