Unconditional Pseudorandomness against Shallow Quantum Circuits

Soumik Ghosh[*1], Sathyawageeswar Subramanian[†2], and Wei Zhan[‡1]

[1]University of Chicago
[2]University of Cambridge

Abstract

Quantum computational pseudorandomness has emerged as a fundamental notion that spans connections to complexity theory, cryptography and fundamental physics. However, all known constructions of efficient quantum-secure pseudorandom objects rely on complexity theoretic assumptions.

In this work, we establish the first *unconditionally* secure efficient pseudorandom constructions against shallow-depth quantum circuit classes. We prove that:

- Any quantum state 2-design yields unconditional pseudorandomness against both QNC^0 circuits with arbitrarily many ancillae and $\mathsf{AC}^0 \circ \mathsf{QNC}^0$ circuits with nearly linear ancillae.

- Random phased subspace states, where the phases are picked using a 4-wise independent function, are unconditionally pseudoentangled against the above circuit classes.

- Any unitary 2-design yields unconditionally secure parallel-query pseudorandom unitaries against geometrically local QNC^0 adversaries, even with limited AC^0 postprocessing.

Our indistinguishability results for 2-designs stand in stark contrast to the standard setting of quantum pseudorandomness against BQP circuits, wherein they can be distinguishable from Haar random ensembles using more than two copies or queries. Our work demonstrates that quantum computational pseudorandomness can be achieved unconditionally for natural classes of restricted adversaries, opening new directions in quantum complexity theory.

---

[*]soumikghosh@uchicago.edu   [†]ss2310@cam.ac.uk   [‡]weizhan@uchicago.edu

1 Introduction

Randomness is fundamental to both classical and quantum computation, particularly in cryptography and algorithm design. However, true randomness is often scarce or computationally impractical. The theory of pseudorandomness studies deterministic objects that appear random to resource-bounded observers. For example, classical pseudorandom generators (PRGs) produce bit strings indistinguishable from random strings for computationally limited observers. A rich theory connects hardness to pseudorandomness for complexity classes such as BPP [Nis91; INW94].

Quantum pseudorandomness extends these ideas to quantum states and unitaries. Truly random quantum objects, described by the Haar measure over the unitary group [Mel24], require exponential resources to generate. Two approaches have emerged for more efficient alternatives. First, the information-theoretic or statistical notion of efficient quantum t-designs: ensembles of states or unitaries that match the first t moments of the Haar measure, and can be prepared or implemented efficiently [BHH16b; BHH16a]. Second, the more recent computational approach: ensembles of quantum pseudorandom states (PRS) or unitaries (PRU) that appear Haar-random to computationally bounded quantum observers [JLS18].

Classically, *unconditional* pseudorandomness has been successfully constructed against several restricted computational models such as constant-depth circuits [Nis91; Bra08; DET+10], read-once branching programs [Nis92; INW94; NZ96], and low-degree polynomials [Vio09; BV10]. These results bypass the need for complexity theoretic assumptions, and have profound implications for cryptography [Lub94], derandomization [SZ99; Hoz21], and lower bounds [HVV06; GMR+12].[1]

In contrast, all existing quantum pseudorandom constructions target powerful adversaries such as polynomial-sized quantum circuits (BQP), and rely on cryptographic assumptions such as the existence of quantum-secure one-way functions.

Our contributions. We establish the first unconditional efficient quantum pseudorandomness results against shallow-depth circuit classes. Such circuits model near-term quantum devices with limited coherence times and gate counts. We show that efficient pseudorandom objects, including PRS, pseudoentanglement, and PRU secure against parallel queries, can be constructed unconditionally for shallow quantum circuits. Our key insight is that due to the depth constraints, each output qubit of shallow quantum circuits locally depends only on a subset of input qubits, thus fundamentally limiting the ability of such circuits to distinguish certain structured quantum objects from Haar-random ones.

A notable aspect of our results is that the only property needed for our constructions is that of being an (approximate) 2-design. *A priori*, the design property only imposes conditions on the behavior of the object when few (in this case exactly two) copies of the objects are present, while pseudorandomness is a property that concerns an arbitrary polynomial number of copies. Rather surprisingly, our results bridge this gap, showing that when the power of the adversaries is restricted, *information-theoretic* indistinguishability on two copies is strong enough to imply *computational* indistinguishability on polynomially many copies.

Our work raises several open questions, ranging from constructing new classes of unconditionally pseudorandom objects against other shallow circuit classes, to applying these results to quantum cryptography and complexity theory. We discuss some of these directions in Section 5, providing new perspectives for analyzing near-term quantum devices.

---

[1] Readers may refer to the survey by Hatami and Hoza [HH24] for a comprehensive review of the more recent developments on classical unconditional pseudorandomness.

1.1 Main results

In this article we construct unconditionally secure efficient pseudorandom objects against two important shallow-depth quantum circuit classes—QNC^0 and $AC^0 \circ QNC^0$.

PRS from state designs. We first demonstrate that unconditional pseudorandomness can be derived from state designs, whose security holds against circuits up to $AC^0 \circ QNC^0$.

Theorem 1.1 (Informal; See Corollary 3.5)**.** Every 2-design state ensemble is an unconditionally secure PRS against QNC^0 circuits with arbitrarily many ancillae and almost linearly many bits of output.

Theorem 1.2 (Informal; See Corollary 3.12)**.** Every 2-design state ensemble is an unconditionally secure PRS against $AC^0 \circ QNC^0$ circuits with almost linearly many ancillae.

Observe that this stands in stark contrast to the case of BQP adversaries. Consider an ensemble that forms at most a t-design, for some $t = \text{poly}(n)$. Then, there are cases—for instance, in the case of random stabilizer states—where a BQP adversary can distinguish between a state from this ensemble and a Haar random state [AD25; GNW21] using more than t copies.

Pseudoentanglement refers to the phenomenon whereby states having very low entanglement are indistinguishable from states having very low entanglement. We also prove that unconditional pseudoentanglement can be achieved against the above shallow quantum circuits.

Theorem 1.3 (Informal; See Corollary 3.13)**.** There exists efficiently constructible, unconditionally secure pseudoentanglement against QNC^0 circuits, and against $AC^0 \circ QNC^0$ circuits with poly-logarithmically many ancillae.

PRU from unitary designs. Similarly, we also prove that unitary t-designs are unconditionally pseudorandom against geometrically local shallow quantum circuits when queried only in parallel (i.e. non-adaptively).

Theorem 1.4 (Informal; See Theorem 4.6)**.** Every unitary 2-design is an unconditionally non-adaptive secure PRU against 1-dimensional geometrically local QNC^0 circuits with arbitrarily many ancillae, and almost linear depth 1-dimensional geometrically local QNC pre-processing.

We also extend this PRU construction to QNC^0 circuits with AC^0 post-processing, with the caveat that it must be weakened slightly since we do not have a natural means to deal with ancillae in the post-processing phase.

Theorem 1.5 (Informal; See Theorem 4.7)**.** Every unitary 2-design on n-qubits is an unconditionally non-adaptive secure PRU against a subclass of $AC^0 \circ QNC^0$ circuits on a multiple of n qubits, and the pre-processing QNC^0 circuit before the queries is 1-dimensional geometrically local.

Proof techniques. A recurring tool that we use in our proofs is that, over any subsystem, the reduced states of a Haar random state are close to maximally mixed with high probability (see Corollary 3.2). Our PRU results require the analogue of this observation for the output of non-adaptive queries to a Haar random unitary (see Corollary 4.4). For this, we bound the expected norm of partial traces of off-diagonal terms $|v\rangle\langle w|$ conjugated by a Haar random unitary (see Lemma 4.3).

To prove our results for QNC^0 with AC^0 post-processing, we observe that when the number of ancillae in the pre-processing QNC^0 circuit is small, the resulting output distribution has high

entropy, although the output distributions are no longer *k*-wise independent. To deal with this, we prove a generalization of Braverman's result [Bra08] that AC^0 circuits cannot distinguish *k*-wise independent distributions from uniform, showing that AC^0 circuits also fail to distinguish *k-wise indistinguishable* distributions with high min-entropy (see Lemma 3.10).

We achieve our pseudoentanglement construction using random phased subspace states, which are superpositions over the orthonormal basis vectors of a subspace with equal amplitudes and random ± 1 phases. We show that such states, when instantiated with a 4-wise independent function for the random phase, are indistinguishable from Haar random by shallow circuits, and have low von Neumann entropy across any cut (see Corollary 3.13).

We believe these technical developments may be of independent interest.

1.2   Related work

Quantum computational notions of pseudorandomness were introduced in [JLS18] and have been studied in a variety of recent works. For instance, many types of pseudoentangled states have been constructed against BQP distinguishers in recent work (for examples, see [ABF+24; BFG+23; BZZ24; ABC+24; GB23; JMW24; FI24]). These notions have found a wide range of applications, from cryptography [AQY22; GE24; AGQ+23] to physics [YE23; GLG+24; FI25; BFG+23; CCG+25]. A number of recent works have also considered the problem of constructing highly efficient pseudorandom unitaries that are implementable in extremely low depth [SHH25; CSB+25].

However, all these constructions rely on complexity theoretic assumptions to obtain pseudorandomness against polynomial-sized quantum circuits. Usually, the assumption has to do with the existence of quantum-secure one way functions [Zha21], based on computational hardness assumptions like the quantum hardness of the learning with errors (LWE) problem [Reg24].

In contrast to this line of work, in our work the adversary is a class of shallow-depth quantum circuits against which we would like our pseudorandom constructions to be secure, and our contribution lies in showing that pseudorandomness against such circuit classes can be obtained without making any complexity theoretic assumptions.

2   Preliminaries

We first define some commonly used notations. We use $[n]$ to denote the set $\{1, \ldots, n\}$. For two distributions $\mathcal{D}$ and $\mathcal{D}'$ over a set X we use $|\mathcal{D} - \mathcal{D}'|_1$ to denote their total variation distance. We denote a random sample x drawn according to $\mathcal{D}$ by $x \sim \mathcal{D}$, and we abuse the notation to denote x drawn uniformly from a set X by $x \sim X$. The identity operator on n qubits is denoted as $\mathbb{I}_n$.

We use $\|\cdot\|_p$ to denote the Schatten-p norms of Hermitian operators. Specifically, $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_\infty$ respectively refers to the trace norm, Frobenius norm and operator norm.

We use the following shorthands for asymptotic growth: $\text{poly}(n) = n^{O(1)}$, $\text{polylog}(n) = \log^{O(1)} n$ and $\text{negl}(n) = n^{-\omega(1)}$.

We assume the readers are familiar with the definitions of the following polynomial-sized circuit classes: QNC for quantum bounded fan-in circuits, AC for classical circuits with unbounded fan-in AND gates, and QAC for quantum circuits with unbounded size CZ gates (but without unbounded fan-out gates). We use QNC^0, AC^0 and QAC^0 to denote their constant-depth subclasses respectively. Without loss of generality, we assume the bounded fan-in is at most 2 (otherwise the constants in some of our results will be changed). For the purpose of this paper, we do not require the quantum circuits to compute cleanly: the ancillae could start with any specified state and also could end up in arbitrary states.

Following [Slo24], we also consider the hybrid circuits with quantum pre-processing and classical post-processing:

Definition 2.1. For a class of classical circuits $\mathcal{F}$ and a class of quantum circuits $\mathcal{C}$, the circuit class $\mathcal{F} \circ \mathcal{C}$ consist of all circuits $F \circ C$ that are composed of a quantum circuit $C \in \mathcal{C}$, followed by computational basis measurements on all output qubits of C, and then with some $F \in \mathcal{F}$ applied on the measurement outcomes. The output distribution of $F \circ C$ with the input state ρ is $F(C(\rho))$.

The class that we are specifically interested in is $\mathsf{AC}^0 \circ \mathsf{QNC}^0$, which is justified in Section 3.2. It is shown in [Slo24] that parity is hard to compute in this class, assuming either no ancillae or linear size of the AC^0 circuit. It worth noting that we do not know yet whether $\mathsf{AC}^0 \circ \mathsf{QNC}^0$ is comparable with QAC^0.

Quantum Pseudorandom Primitives

Below we generalize the commonly used definitions of quantum pseudorandom primitives to those with respect to specific classes of adversaries, rather than simply polynomial-time adversaries. The state and unitary ensembles are all discrete distributions, which we denote by their supports for succinctness.

Definition 2.2. The state ensemble $\{|\psi_i\rangle\}$ on n qubits is a *pseudorandom state ensemble (PRS)* against a class of quantum circuits $\mathcal{C}$, if for the n-qubit Haar random state $|\psi_{\mathsf{Haar}}\rangle$, every $t = \mathrm{poly}(n)$ and every circuit $C \in \mathcal{C}$, we have

$$\left| \mathbb{E}_i \left[C \left(|\psi_i\rangle\langle\psi_i|^{\otimes t} \right) \right] - \mathbb{E} \left[C \left(|\psi_{\mathsf{Haar}}\rangle\langle\psi_{\mathsf{Haar}}|^{\otimes t} \right) \right] \right|_1 = \mathrm{negl}(n).$$

Here $C(\rho)$ represents the output distribution with input state ρ.

Definition 2.3. We say two state ensembles $\{|\psi_i\rangle\}$ and $\{|\psi_j\rangle\}$ on n qubits demonstrate *pseudoentanglement* against a class of quantum circuits $\mathcal{C}$, if every for every $t = \mathrm{poly}(n)$ and every circuit $C \in \mathcal{C}$, we have

$$\left| \mathbb{E}_i \left[C \left(|\psi_i\rangle\langle\psi_i|^{\otimes t} \right) \right] - \mathbb{E}_j \left[C \left(|\psi_j'\rangle\langle\psi_j'|^{\otimes t} \right) \right] \right|_1 = \mathrm{negl}(n),$$

while across the same bipartition or cut of the qubits, the expected entanglement entropies of $\{|\psi_i\rangle\}$ and $\{|\psi_j'\rangle\}$ are asymptotically different.

Definition 2.4. The unitary ensemble $\{U_i\}$ on n qubits is a *pseudorandom unitary ensemble (PRU)* against a class of quantum circuits $\mathcal{C}$, if for the n-qubit Haar random unitary U_{Haar} and every circuit $C^U \in \mathcal{C}^U$ on $t = \mathrm{poly}(n)$ input qubits, we have

$$\left| \mathbb{E}_i \left[C^{U_i}(|0^t\rangle\langle 0^t|) \right] - \mathbb{E} \left[C^{U_{\mathsf{Haar}}}(|0^t\rangle\langle 0^t|) \right] \right|_1 = \mathrm{negl}(n). \tag{1}$$

Here C^U stands for a circuit C which uses U as oracle gates.

In this work we are concerned with the notion of PRU when U is guaranteed to be applied in parallel, that is, (1) is only required to hold for circuits C^U that apply all their U gates in a single layer. In this case, we say $\{U_i\}$ on n qubits is a *parallel-query* (or *non-adaptive-query*, as defined in [MPS+24]) *secure PRU* against $\mathcal{C}$. We refer to the part of the circuit C^U before the layer of U gates as pre-processing, and the part after as post-processing.

We will also discuss some potential constructions of unconditional pseudorandom generators against shallow quantum circuits, defined as follows.

Definition 2.5. The boolean function $G : \{0,1\}^\ell \to \{0,1\}^n$ is a *t-copy pseudorandom generator (PRG)* against a class of quantum circuits $\mathcal{C}$, if for every circuit $C \in \mathcal{C}$, we have

$$\left| \mathop{\mathbb{E}}_{x \sim \{0,1\}^\ell} \left[C \left(|G(x)\rangle \langle G(x)|^{\otimes t} \right) \right] - \mathop{\mathbb{E}}_{x \sim \{0,1\}^n} \left[C \left(|x\rangle \langle x|^{\otimes t} \right) \right] \right|_1 = \mathrm{negl}(n).$$

In particular, G is a *pseudorandom generator* against $\mathcal{C}$ if it is a t-copy pseudorandom generator for every $t = \mathrm{poly}(n)$.

State and Unitary Designs

Here we recall the definitions and properties of exact and approximate designs, which are statistical notions of pseudorandomness.

Definition 2.6 (State t-design). The state ensemble $\{|\psi_i\rangle\}$ on n qubits is a *t-design*, if for the n-qubit Haar random unitary $|\psi_{\mathsf{Haar}}\rangle$, we have

$$\mathop{\mathbb{E}}_i \left[|\psi_i\rangle \langle \psi_i|^{\otimes t} \right] = \mathbb{E} \left[|\psi_{\mathsf{Haar}}\rangle \langle \psi_{\mathsf{Haar}}|^{\otimes t} \right].$$

We say that $\{|\psi_i\rangle\}$ is an *ε-approximate t-design*, if instead we have

$$\left\| \mathop{\mathbb{E}}_i \left[|\psi_i\rangle \langle \psi_i|^{\otimes t} \right] - \mathbb{E} \left[|\psi_{\mathsf{Haar}}\rangle \langle \psi_{\mathsf{Haar}}|^{\otimes t} \right] \right\|_1 \leq \varepsilon.$$

Lemma 2.7. Let $\{|\psi_i\rangle\}$ be an ϵ-approximate state 2-design on n qubits, and let B be a subsystem of containing $n - k$ qubits. We have

$$\mathop{\mathbb{E}}_i \left[\| \mathrm{Tr}_B[|\psi_i\rangle \langle \psi_i|] \|_2^2 \right] \leq \mathbb{E} \left[\| \mathrm{Tr}_B[|\psi_{\mathsf{Haar}}\rangle \langle \psi_{\mathsf{Haar}}|] \|_2^2 \right] + \varepsilon.$$

Proof. Denote the complementary subsystem to B by A, which contains k qubits. Define the swap operator R by the identity

$$\mathrm{Tr}[\mathrm{Tr}_B \rho_1 \cdot \mathrm{Tr}_B \rho_2] = \mathrm{Tr}\left[(\rho_1 \otimes \rho_2) \cdot R \right],$$

where one can check that

$$R = \sum_{x,y \in \{0,1\}^k} |x\rangle \langle y|_A |y\rangle \langle x|_{A'} \otimes \mathbb{I}_{BB'}.$$

Here A' and B' are identical copies of A and B respectively.

Since R is a permutation matrix, the operator norm of R is exactly 1. As a result, by Hölder's inequality we have

$$\mathop{\mathbb{E}}_i \left[\| \mathrm{Tr}_B[|\psi_i\rangle \langle \psi_i|] \|_2^2 \right] - \mathbb{E} \left[\| \mathrm{Tr}_B[|\psi_{\mathsf{Haar}}\rangle \langle \psi_{\mathsf{Haar}}|] \|_2^2 \right]$$

$$= \mathrm{Tr}\left[\left(\mathop{\mathbb{E}}_i \left[|\psi_i\rangle \langle \psi_i|^{\otimes 2} \right] - \mathbb{E} \left[|\psi_{\mathsf{Haar}}\rangle \langle \psi_{\mathsf{Haar}}|^{\otimes 2} \right] \right) \cdot R \right]$$

$$= \left\| \mathop{\mathbb{E}}_i \left[|\psi_i\rangle \langle \psi_i|^{\otimes t} \right] - \mathbb{E} \left[|\psi_{\mathsf{Haar}}\rangle \langle \psi_{\mathsf{Haar}}|^{\otimes t} \right] \right\|_1 \cdot \|R\|_\infty \leq \varepsilon. \qquad \square$$

Definition 2.8 (Unitary t-design). Let $\mathcal{D} = \{U_i\}$ be an ensemble of n-qubit unitaries. The unitary ensemble $\mathcal{D}$ is a *unitary t-design* if, for the n-qubit Haar-random unitary U_{Haar}, we have

$$\mathbb{E}_i \left[U_i^{\otimes t} \otimes (U_i^\dagger)^{\otimes t} \right] = \mathbb{E} \left[U_{\text{Haar}}^{\otimes t} \otimes (U_{\text{Haar}}^\dagger)^{\otimes t} \right].$$

Define the t-th moment channels as

$$\Phi_{\mathcal{D}}^{(t)}(\rho) = \mathbb{E}_i \left[U_i^{\otimes t} \cdot \rho \cdot (U_i^\dagger)^{\otimes t} \right], \qquad \Phi_{\text{Haar}}^{(t)}(\rho) = \mathbb{E} \left[U_{\text{Haar}}^{\otimes t} \cdot \rho \cdot (U_{\text{Haar}}^\dagger)^{\otimes t} \right],$$

we say that $\mathcal{D}$ is an *ε-approximate unitary t-design*, if for all operators ρ with $\|\rho\|_1 \leq 1$ we have

$$\left\| \Phi_{\mathcal{D}}^{(t)}(\rho) - \Phi_{\text{Haar}}^{(t)}(\rho) \right\|_1 \leq \varepsilon.$$

Lemma 2.9. Let $\{U_i\}$ be an ϵ-approximate unitary 2-design on n qubits, and let B be a subsystem of the n qubits. For every two n-qubit states $|v\rangle$ and $|w\rangle$, we have

$$\mathbb{E}_i \left[\left\| \text{Tr}_B[U_i|v\rangle\langle w|U_i^\dagger] \right\|_2^2 \right] \leq \mathbb{E} \left[\left\| \text{Tr}_B[U_{\text{Haar}}|v\rangle\langle w|U_{\text{Haar}}^\dagger] \right\|_2^2 \right] + \varepsilon.$$

Proof. Denote the complementary subsystem to B by A, and assume that A contains k qubits. Define the operator R the same way as the proof above for Lemma 2.7, and we have

$$\left\| \text{Tr}_B(U|v\rangle\langle w|U^\dagger) \right\|_2^2 = \text{Tr}\left[(U|v\rangle\langle w|U^\dagger)^{\otimes 2} \cdot R \right].$$

Similarly, since $\|R\|_\infty = 1$ and $\left\| |v\rangle\langle w|^{\otimes 2} \right\|_1 = 1$, by Hölder's inequality we have

$$\mathbb{E}_i \left[\left\| \text{Tr}_B[U_i|v\rangle\langle w|U_i^\dagger] \right\|_2^2 \right] - \mathbb{E} \left[\left\| \text{Tr}_B[U_{\text{Haar}}|v\rangle\langle w|U_{\text{Haar}}^\dagger] \right\|_2^2 \right]$$

$$= \text{Tr}\left[\left(\Phi_{\mathcal{D}}^{(t)}(|v\rangle\langle w|^{\otimes 2}) - \Phi_{\text{Haar}}^{(t)}(|v\rangle\langle w|^{\otimes 2}) \right) \cdot R \right]$$

$$= \left\| \Phi_{\mathcal{D}}^{(t)}(|v\rangle\langle w|^{\otimes 2}) - \Phi_{\text{Haar}}^{(t)}(|v\rangle\langle w|^{\otimes 2}) \right\|_1 \cdot \|R\|_\infty \leq \varepsilon. \qquad \square$$

Schmidt Decomposition

We will need several facts about the Schmidt decomposition (listed below), whose proofs can be found in e.g. [NC10].

Definition 2.10. Let $\mathcal{H}_1, \mathcal{H}_2$ be two Hilbert spaces, and let $x \in \mathcal{H}_1 \otimes \mathcal{H}_2$. If we write

$$x = \sum_{i=1}^{r} \alpha_i \cdot v_i \otimes w_i, \tag{2}$$

where $\alpha_i \in \mathbb{C}$, $v_i \in \mathcal{H}_1$ and $w_i \in \mathcal{H}_2$, we call (2) a *tensor product decomposition* of x. Furthermore, if both $\{v_i\}$ and $\{w_i\}$ are orthonormal and each α_i is non-zero, we call (2) a *Schmidt decomposition* and r the *Schmidt rank* of x with respect to $\mathcal{H}_1$ and $\mathcal{H}_2$.

Lemma 2.11. Let $\mathcal{H}_1, \mathcal{H}_2$ be two Hilbert spaces, and let $x \in \mathcal{H}_1 \otimes \mathcal{H}_2$. Then:

- In any tensor product decomposition of x as in (2), the number of terms r is at least the Schmidt rank of x;

- Let $\{v_i\}$ and $\{w_j\}$ be two orthonormal basis for $\mathcal{H}_1$ and $\mathcal{H}_2$, respectively. If we write

$$x = \sum_{i,j} \alpha_{ij} \cdot v_i \otimes w_j,$$

  then the Schmidt rank of x is exactly the rank of the matrix with entry α_{ij} at the i-th row and j-th column.

- The von Neumann entanglement entropy of x, with respect to the subsystems $\mathcal{H}_1$ and $\mathcal{H}_2$, is at most $\log_2 r$ where r is the Schmidt rank of x.

3 Unconditional pseudorandomness from 2-designs

In this section, we focus on state designs which exploit the locality properties of shallow circuits in order to achieve unconditional pseudorandomness. At a high level, this resembles a quantum analog of small bias distributions (e.g. [NN93]), which can fool low-degree polynomials. We will begin with some facts about Haar random states, which relate the size of the subsystem with entanglement entropy, and allow us to approximate small subsystems with maximally mixed states.

Lemma 3.1 ([Lub78; LLZ+18]). Let $|\psi_{\mathsf{Haar}}\rangle$ be an n-qubit Haar random state, and let ρ_A be the reduced density matrix of $|\psi_{\mathsf{Haar}}\rangle\langle\psi_{\mathsf{Haar}}|$ on the subsystem A with $|A| = k$ qubits. Then

$$\mathbb{E}[\operatorname{Tr}(\rho_A^2)] = \frac{2^k + 2^{n-k}}{2^n + 1}.$$

Corollary 3.2. Let $|\psi_{\mathsf{Haar}}\rangle$ be an n-qubit Haar random state. For any $t \geq 1$, let ρ_A be the reduced density matrix of $|\psi_{\mathsf{Haar}}\rangle\langle\psi_{\mathsf{Haar}}|^{\otimes t}$ over a subsystem A. Then for every $\delta > 0$, with probability at least $1 - n^{O(k)} \cdot 2^{-n/2} \cdot \delta^{-1}$ over $|\psi_{\mathsf{Haar}}\rangle$, it holds for all A with $|A| = k$ qubits that

$$\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_1 \leq \delta.$$

Proof. First consider the case when A is fully contained in one copy of the Haar random state. In this case from Lemma 3.1 we have

$$
\begin{aligned}
\mathbb{E}\left[\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_1\right] &\leq 2^{k/2} \cdot \mathbb{E}\left[\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_2\right] \\
&\leq 2^{k/2} \cdot \mathbb{E}\left[\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_2^2\right]^{1/2} \\
&= 2^{k/2} \cdot \mathbb{E}\left[\operatorname{Tr}(\rho_A^2) - 2^{-k}\right]^{1/2} \\
&= 2^{k/2} \cdot \left(\frac{2^k - 2^{-k}}{2^n + 1}\right)^{1/2} \\
&\leq 2^{k-n/2}.
\end{aligned}
$$

By Markov's inequality, we know that $\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_1 \leq \delta/k$ holds with probability at least $1 - k \cdot 2^{k-n/2} \cdot \delta^{-1}$. By a union bound, this holds for all A with $|A| \leq k$ with probability at least $1 - n^{O(k)} \cdot 2^{-n/2} \cdot \delta^{-1}$.

When A consists qubits from at most k different copies, we denote the subsystems as $A = A_1 \sqcup A_2 \sqcup \cdots$ with $|A_i| = k_i$. Since the copies are unentangled with each other, we have $\rho_A = \rho_{A_1} \otimes \rho_{A_2} \otimes \cdots$, and thus

$$\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_1 \leq \sum_i \left\| \rho_{A_1} \otimes \cdots \otimes \rho_{A_i} - \rho_{A_1} \otimes \cdots \otimes \rho_{A_{i-1}} \otimes \frac{1}{2^{k_i}} \mathbb{I}_{k_i} \right\|_1$$

$$= \sum_i \left\| \rho_{A_i} - \frac{1}{2^{k_i}} \mathbb{I}_{k_i} \right\|_1$$

$$\leq \delta$$

with probability at least $1 - n^{O(k)} \cdot 2^{-n/2} \cdot \delta^{-1}$. $\qquad\square$

Notice that the proof of Corollary 3.2 only uses the second moment properties of $|\psi_{\mathsf{Haar}}\rangle$, and therefore the conclusions immediately hold for approximate 2-designs with negligible error as well.

Corollary 3.3. Let $\{|\psi_i\rangle\}$ be an ε-approximate 2-design on n qubits. For any $t \geq 1$, let ρ_A be the reduced density matrix of $|\psi_i\rangle\langle\psi_i|^{\otimes t}$ over a subsystem A. Then for every $\delta > 0$, with probability at least $1 - n^{O(k)} \cdot (\varepsilon + 2^{-n})^{1/2} \cdot \delta^{-1}$ over i, it holds for all A with $|A| = k$ qubits that

$$\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_1 \leq \delta.$$

Proof. By Lemma 2.7, the approximate design property implies that

$$\mathbb{E}\left[\mathrm{Tr}(\rho_A^2) - 2^{-k} \right]^{1/2} \leq (\varepsilon + 2^{k-n})^{1/2},$$

and thus

$$\mathbb{E}\left[\left\| \rho_A - \frac{1}{2^k} \mathbb{I}_k \right\|_1 \right] \leq 2^k \cdot (\varepsilon + 2^{-n})^{1/2}.$$

The rest of the proof follows the same arguments from Corollary 3.2. $\qquad\square$

3.1 Pseudorandomness against QNC^0

As a warm-up, we will use Corollary 3.2 and Corollary 3.3 to show that any 2-design is indistinguishable to a Haar random state, with respect to any QNC^0 distinguisher.

Theorem 3.4. Let $\{|\psi_i\rangle\}$ be an ε-approximate 2-design on n qubits for some $\varepsilon = \mathrm{negl}(n)$. Then $\{|\psi_i\rangle\}$ is a PRS against QNC circuits with depth

$$d = \min(\log\log(1/\varepsilon), \log n) - \log\log n - \omega(1)$$

and a single-bit output. In particular, when $\varepsilon \leq 2^{-\Omega(n)}$, $\{|\psi_i\rangle\}$ is PRS against QNC circuits up to depth $d = \log n - \log\log n - \omega(1)$.

Proof. The output bit of the depth-d QNC circuit depends on at most $k = 2^d$ input qubits and ancillae. Let A be the subsystem of the input state on these qubits, and let m be the number of ancillae touched. Denote the reduced density matrix, over the subsystem A, of the Haar random

9

state to be ρ_A^{Haar} and that of the state picked from $\{|\psi_i\rangle\}$ to be ρ_A. Then for the channel Φ_C that maps the subsystem A to the output qubit, we have

$$\left\|\Phi_C\left(\rho_A^{\mathsf{Haar}}\right) - \Phi_C\left(\rho_A\right)\right\|_1 \le \left\|\rho_A^{\mathsf{Haar}} - \rho_A\right\|_1 = \mathrm{negl}(n),$$

with probability at least $1 - \mathrm{negl}(n)$ over the choice of the states. This follows from Corollary 3.2 and Corollary 3.3, and the observation that

$$n^{O(k)} \cdot (\varepsilon + 2^{-n})^{1/2} \le \mathrm{negl}(n)$$

is equivalent to

$$O(k) \le \frac{\log(1/(\varepsilon + 2^{-n}))}{2\log n} - \omega(1),$$

which is satisfied for every $d = \log k$ such that

$$d \le \min(\log\log(1/\varepsilon), \log n) - \log\log n - \omega(1). \qquad \square$$

The above theorem can be strengthened to work for the case when multiple output qubits are measured.

Corollary 3.5. Let $\{|\psi_i\rangle\}$ be an ε-approximate 2-design on n qubits for some $\varepsilon = 2^{-\Omega(n)}$. Then $|\psi_i\rangle$ is a PRS against QNC circuits with depth d and $k \le 2^{-d} \cdot o(n/\log n)$ bits of output.

Proof. The backward lightcone of the k output qubits is of at most

$$k \cdot 2^d = o(n/\log n)$$

in size. The proof then follows from the same argument as Theorem 3.4. $\qquad \square$

3.2 Pseudorandomness against $\mathsf{AC}^0 \circ \mathsf{QNC}^0$

The lightcone argument in the previous section renders most parts of a QNC^0 circuit and most input qubits unrelated. Naturally, it is more desirable to prove the statement against QNC^0 circuits where all the output qubits are measured, that is, the output distribution is indistinguishable in total variation distance between t-designs and Haar random states. However, our example below shows that this is too much to ask for, even when the QNC^0 circuit does nothing and the input states get measured immediately.

Definition 3.6 (Random phased subspace states). A d-dimensional *random phased subspace state* on n qubits is the following state:

$$|\psi_{S,f}\rangle = \frac{1}{2^{d/2}} \sum_{x \in S} (-1)^{f(x)} |x\rangle,$$

where $S \in \mathbb{F}_2^n$ is a random d-dimensional linear subspace, and $f : S \to \{0,1\}$ is a random function.

Theorem 3.7. For every $n > d > t$, $\{|\psi_{S,f}\rangle\}$ is an $O(2^{t-d})$-approximate t-design, but $|\psi_{S,f}\rangle^{\otimes(d+1)}$ and $|\psi_{\mathsf{Haar}}\rangle^{\otimes(d+1)}$ are distinguishable when measured in the computational basis, with total variation distance $1 - O(2^d)/2^n$.

Proof. The proof that $\{|\psi_{S,f}\rangle\}$ is an approximate design is deferred to Appendix A. To distinguish $d+1$ copies of $\{|\psi_{S,f}\rangle\}$ from Haar random, notice that the measurement outcome in each copy is a random $x \in S$, and thus the $d+1$ outcomes must be linearly dependent. On the other hand, the measurement outcomes from $d+1$ copies of a Haar random state, when all distinct, form a random $(d+1)$-element subset of $\{0,1\}^n$ because of symmetry, and therefore are linearly dependent with probability at most

$$\frac{1}{2^n} + \frac{1}{2^{n-1}} + \cdots + \frac{1}{2^{n-d}} \le \frac{1}{2^{n-d-1}}.$$

We also know that the collision probability for a Haar random state is $2/(2^n + 1)$ [DHB22], and thus the probability for the outcomes not being all distinct is at most $d(d+1)/2^n$ by the union bound. As a result, the total variation distance between the measurement outcomes of $|\psi_{S,f}\rangle^{\otimes(d+1)}$ and $|\psi_{\mathsf{Haar}}\rangle^{\otimes(d+1)}$ is at least $1 - 1/2^{n-d-1} - d(d+1)/2^n = 1 - O(2^d)/2^n$. $\square$

Notice that not only the distinguisher in Theorem 3.7 applies no quantum gates, the classical post-processing on the measurement outcomes are also quite simple. It checks linear dependence whose complexity is captured by DET, a complexity class between NC^1 and NC^2 containing all problems reducible to determinant. This motivates us to examine the case when the classical post-processing is restricted to some provably weaker class than DET. It turns out that for AC^0, 2-designs are indeed still pseudorandom in this case. We crucially make use of the follow result, first proved by Braverman [Bra08], subsequentially improved by [Tal17] and Harsha and Srinivasan [HS19], that almost k-wise uniform distributions fools AC^0:

Lemma 3.8. *For every δ-almost k-wise independent distribution on m bits, any AC circuits with size s and depth d cannot distinguish it from the uniform distribution with advantage $\varepsilon + m^k\delta$, for certain $k = (\log s)^{O(d)} \cdot \log(1/\varepsilon)$.*

We start out with the simple case, when no ancilla is allowed for the QNC^0 circuit.

Theorem 3.9. *Let $\{|\psi_i\rangle\}$ be an ε-approximate 2-design on n qubits for some $\varepsilon = 2^{-\log^{\omega(1)} n}$. Then $|\psi_i\rangle$ is a PRS against $\mathsf{AC}^0 \circ \mathsf{QNC}^0$ circuits with no ancilla.*

Proof. Suppose the circuit has size s and depth d. Fix some $k = (\log s)^{O(d)} \cdot \log^2 n = \log^{O(1)} n$ according to Lemma 3.8.

The output of the QNC^0 circuit, when all qubits are measured, is a distribution $\mathcal{D}_{|\psi\rangle}$ over $m = \mathrm{poly}(n)$ bits that depends on the input state $|\psi\rangle$. By Corollary 3.3, for both $\mathcal{D}_{|\psi_i\rangle}$ and $\mathcal{D}_{|\psi_{\mathsf{Haar}}\rangle}$, with probability $1 - \mathrm{negl}(n)$, the marginal distribution on every k bits are δ-indistinguishable from the case when the input states are maximally mixed, for every $\delta > 0$ such that $n^{O(k)} \cdot (\varepsilon + 2^{-n})^{1/2} \cdot \delta^{-1} = \mathrm{negl}(n)$.

Since there is no ancilla, the output distribution when the inputs are maximally mixed is the uniform distribution. Therefore both $\mathcal{D}_{|\psi_i\rangle}$ and $\mathcal{D}_{|\psi_{\mathsf{Haar}}\rangle}$ are δ-almost k-wise independent. By Lemma 3.8 both distributions are indistinguishable from the uniform distribution against the AC^0 post-processing, as long as $m^k\delta = n^{O(k)}\delta$ is negligible. This is satisfied by our choice of ε, as

$$n^{O(k)} \cdot (\varepsilon + 2^{-n})^{1/2} = 2^{\log^{O(1)} n - \log^{\omega(1)} n} = \mathrm{negl}(n). \qquad \square$$

Notice that for exact 2-designs, that is when $\varepsilon = 0$, we only need $k = o(n/\log n)$, and thus Theorem 3.9 can be strengthened to work against $\mathsf{AC} \circ \mathsf{QNC}$ circuits of polynomial size, QNC depth up to $o(\log n)$ and AC depth up to $o(\log n/\log\log n)$.

The situation becomes more complicated when ancillae are allowed. In this case, although $\mathcal{D}_{|\psi_i\rangle}$ and $\mathcal{D}_{|\psi_{\mathsf{Haar}}\rangle}$ are still almost k-wise indistinguishable (that the marginal distribution on every k bits

are close in total variation distance), they are no longer k-wise independent and are not guaranteed to fool AC^0 circuits when k is small [BIV$^+$16]. This happens because we have no knowledge of the output distribution of the QNC^0 circuit with ancillae even when the inputs are maximally mixed. The saving grace is that, when the number of ancillae is small, the output distribution has high entropy and we can modify Braverman's proof in [Bra08] to suit such distributions:

Lemma 3.10. For every two δ-almost k-wise indistinguishable distributions $\mathcal{D}_1, \mathcal{D}_2$ on m bits, such that $\mathcal{D}_1$ has min-entropy at least $m - r$, any AC circuits with size s and depth d cannot distinguish the two distributions with advantage $\varepsilon + 4m^k\delta$, for certain $k = (\log s)^{O(d)} \cdot (r + \log(1/\varepsilon))$.

Proof. We first consider the case when $\delta = 0$. Suppose the AC^0 circuit with size s and depth d computes a boolean function F, Braverman showed that [Bra08, Lemma 11] there exists a boolean function F' and polynomial f' of degree $k = (\log s)^{O(d)} \cdot \log(1/\varepsilon')$, such that:

- $\mathbb{P}_{\mathcal{D}_2}[F \neq F'] < \varepsilon'$,

- $\mathbb{P}_{\mathcal{U}}[F \neq F'] < \varepsilon'$,

- $F' \geq f'$ on $\{0,1\}^m$ and $\mathbb{E}_{\mathcal{U}}[F' - f'] < \varepsilon'$.

Here $\mathcal{U}$ stands for the uniform distribution over $\{0,1\}^m$. Since $\mathcal{D}_1$ has min-entropy at least $m - r$, we have $\mathbb{P}_{\mathcal{D}_1}[F \neq F'] < 2^r\varepsilon$ and $\mathbb{E}_{\mathcal{D}_1}[F' - f'] < 2^r\varepsilon$. As a result,

$$
\begin{aligned}
\mathbb{E}_{\mathcal{D}_2}[F] > \mathbb{E}_{\mathcal{D}_2}[F'] - \varepsilon' &\geq \mathbb{E}_{\mathcal{D}_2}[f'] - \varepsilon' \\
&= \mathbb{E}_{\mathcal{D}_1}[f'] - \varepsilon' > \mathbb{E}_{\mathcal{D}_1}[F'] - (2^r + 1)\varepsilon' \\
&> \mathbb{E}_{\mathcal{D}_1}[F] - (2^{r+1} + 1)\varepsilon'.
\end{aligned}
$$

The bound in the reverse direction also holds by considering $1 - F$. Taking $\varepsilon' = \varepsilon/(2^{r+1} + 1)$ gives us $|\mathbb{E}_{\mathcal{D}_1}[F] - \mathbb{E}_{\mathcal{D}_2}[F]| < \varepsilon$.

For $\delta > 0$, we can assume that $m^k\delta < 1$, as otherwise the lemma is trivial. By [BIV$^+$16], there exists two k-wise indistinguishable distributions $\mathcal{D}'_1, \mathcal{D}'_2$ on m bits such that $|\mathcal{D}_1 - \mathcal{D}'_1|_1 \leq 2m^k\delta$ and $|\mathcal{D}_2 - \mathcal{D}'_2|_1 \leq 2m^k\delta$. Furthermore, the construction ensures that $\|\mathcal{D}'_1\|_\infty \leq \|\mathcal{D}_1\|_\infty + 2m^k\delta \cdot 2^{-m}$, and thus $\mathcal{D}'_1$ still has min-entropy at least $m - O(r)$. Applying the previous claim on $\mathcal{D}'_1, \mathcal{D}'_2$ we have $\left|\mathbb{E}_{\mathcal{D}'_1}[F] - \mathbb{E}_{\mathcal{D}'_2}[F]\right| < \varepsilon$, and thus $|\mathbb{E}_{\mathcal{D}_1}[F] - \mathbb{E}_{\mathcal{D}_2}[F]| < \varepsilon + 4m^k\delta$. $\qquad\square$

Theorem 3.11. Let $\{|\psi_i\rangle\}$ be an ε-approximate 2-design on n qubits for some $\varepsilon = 2^{-\log^{\omega(1)} n}$. Then $|\psi_i\rangle$ is a PRS against $\mathsf{AC}^0 \circ \mathsf{QNC}^0$ circuits with $a = \mathrm{polylog}(n)$ ancillae.

Proof. The forward lightcone of the ancillae touches at most $r = 2^d a = \log^{O(1)} n$ output qubits. We called the these qubits *corrupted*, denoted as the subsystem R. We define $\mathcal{R}$ as the distribution over $\{0,1\}^r$ following the measurement outcome on the corrupted qubits when the input states are maximally mixed.

Similar to the proof of Theorem 3.9, suppose the circuit has size s and depth d, and we fix some $k = (\log s)^{O(d)} \cdot (r + \log^2 n) = \log^{O(1)} n$ according to Lemma 3.10. Denote the output distribution of the QNC^0 circuit as $\mathcal{D}_{|\psi\rangle}$ with the input state $|\psi\rangle$. Let $\delta > 0$ satisfies $n^{O(k)} \cdot (\varepsilon + 2^{-n})^{1/2} \cdot \delta^{-1} = \mathrm{negl}(n)$. We claim that with probability $1 - \mathrm{negl}(n)$, both $\mathcal{D}_{|\psi_i\rangle}$ and $\mathcal{D}_{|\psi_{\mathsf{Haar}}\rangle}$ are δ-almost k-wise indistinguishable from $\mathcal{U} \otimes \mathcal{R}$, where $\mathcal{U}$ is the uniform distribution over the uncorrupted output bits.

Figure 1: An illustration of the proof of Theorem 3.11, with partial systems and unitary lightcones in an QNC^0 circuit.

To prove the claim, let A be the ancilla qubits, and ρ_A be the initial state of the ancillae. Let the unitary operator U_R consist of all gates in the QNC^0 circuit that belongs to the lightcone of the ancillae. Fixing any k output qubits K, we extending K to $K \cup R$ so that it contains all the corrupted qubits, and denote U_B as the unitary operator that consists of all gates in the backward lightcone of K but outside U_R. Notice that we can think of U_R as being applied after U_B. At the input, the backward lightcone of K touches the set of non-ancilla qubits B with $|B| = b$, and we let ρ_B be the state of B when the input states are copies of $|\psi\rangle$.

Now the output state on K is a partial trace of

$$(\mathbb{I}_{A \cup B \setminus R} \otimes U_R)(U_B \rho_B U_B^\dagger \otimes \rho_A)(\mathbb{I}_{A \cup B \setminus R} \otimes U_R^\dagger), \tag{3}$$

and it suffices to show that no matter if $|\phi\rangle$ is drawn from $\{|\psi_i\rangle\}$ or $|\psi_{\text{Haar}}\rangle$, with high probability the above state is close to the case when ρ_B is maximally mixed. Indeed, by Corollary 3.2 and Corollary 3.3, in both former cases with probability $1 - \text{negl}(n)$ we have $\left\| \rho_B - 2^{-(k-a)} \mathbb{I}_B \right\|_1 \leq \delta$. This means that the state in (3) if δ-close in trace distance to

$$(\mathbb{I}_{A \cup B \setminus R} \otimes U_R)(2^{-b} \mathbb{I}_B \otimes \rho_A)(\mathbb{I}_{A \cup B \setminus R} \otimes U_R^\dagger) = 2^{-(a+b-r)} \mathbb{I}_{A \cup B \setminus R} \otimes 2^{-(r-a)} U_R(\mathbb{I}_{R \setminus A} \otimes \rho_A) U_R^\dagger.$$

Notice that $2^{-(r-a)} U_R(\mathbb{I}_{R \setminus A} \otimes \rho_A) U_R^\dagger$ is exactly the output state on R when the input states are maximally mixed, and therefore the measurement outcome has the distribution $\mathcal{R}$. Meanwhile $2^{-(a+b-r)} \mathbb{I}_{A \cup B \setminus R}$ is maximally mixed and will be measured to the uniform distribution. Thus we conclude that the output distributions on the k bits are δ-close to $\mathcal{U} \otimes \mathcal{R}$, for both $\mathcal{D}_{|\psi_i\rangle}$ and $\mathcal{D}_{|\psi_{\text{Haar}}\rangle}$.

Now we use the fact that $\mathcal{U} \otimes \mathcal{R}$, as a distribution over $\{0,1\}^m$, has min-entropy at least $m - r$. By Lemma 3.10, both $\mathcal{D}_{|\psi_i\rangle}$ and $\mathcal{D}_{|\psi_{\text{Haar}}\rangle}$ are indistinguishable from $\mathcal{U} \otimes \mathcal{R}$ against the AC^0 post-processing, as long as $m^k \delta = n^{O(k)} \delta$ is negligible. Similar to the proof of Theorem 3.9, this is satisfied by our choice of ε. $\qquad \square$

Similar to the case of Theorem 3.9, for exact 2-designs, Theorem 3.11 can be strengthened to work against AC $\circ$ QNC circuits of polynomial size, QNC depth up to $o(\log n)$ and AC depth up to $o(\log n / \log \log n)$. On the other hand, for constant depth circuits the number of ancillae can be strengthened close to linear.

Corollary 3.12. Let $\{|\psi_i\rangle\}$ be an exact 2-design on n qubits. Then $|\psi_i\rangle$ is a PRS against $\mathsf{AC}^0 \circ \mathsf{QNC}^0$ circuits with $a = n/\log^{\omega(1)} n$ ancillae.

Proof. In the proof of Theorem 3.11, when $\varepsilon = 0$ we can take $\delta = 2^{-n/4}$, and thus to have $m^k\delta = \mathrm{negl}(n)$ for any $m = \mathrm{poly}(n)$ we only need $k = o(n/\log n)$. When $d = O(1)$, this is satisfies by any $a = 2^{-d}r = n/\log^{\omega(1)} n$. □

3.3 Pseudoentanglement against $\mathsf{AC}^0 \circ \mathsf{QNC}^0$

In contrast to prior works that constructed pseudoentanglement from quantum secure one-way functions, here we prove that unconditional pseudoentanglement is possible against shallow circuits, using random phased subspace states (see Definition 3.6). We will show that such states form good enough approximate t-designs, even when the phases are picked using a $2t$-wise independent function, and thus yield pseudorandomness by our previous results.

Corollary 3.13. Let $\{|\psi_{S,f}\rangle\}$ be the ensemble of d-dimensional random phased subspace states, instantiated with a 4-wise independent function $f : \{0,1\}^n \to \{0,1\}$. Then the following properties hold:

- The ensemble is indistinguishable from a Haar random ensemble against:
    - QNC^0 circuits, when $d = \omega(\log n)$;
    - $\mathsf{AC}^0 \circ \mathsf{QNC}^0$ circuits with $\mathrm{polylog}(n)$ ancillae, when $d = \log^{\omega(1)} n$;
- The von Neumann entanglement entropy across any cut is at most d.

Proof. First notice that even when f is 4-wise independent (or in general, $2t$-wise independent) instead of truly random, the proof of the design property of $\{|\psi_{S,f}\rangle\}$ in Appendix A still holds. Specifically, in Equation (9):

$$\mathop{\mathbb{E}}_f \left[|\psi_{S,f}\rangle\langle\psi_{S,f}|^{\otimes t}\right] = \frac{1}{2^{dt}} \sum_{\substack{x_1,\ldots,x_t \in S \\ y_1,\ldots,y_t \in S}} \mathop{\mathbb{E}}_f \left[(-1)^{f(x_1)+\cdots+f(x_t)+f(y_1)+\cdots+f(y_t)}\right] |x_1 \cdots x_t\rangle\langle y_1 \cdots y_t|,$$

When f is $2t$-wise independent, the expectation of $(-1)^{f(x_1)+\cdots+f(x_t)+f(y_1)+\cdots+f(y_t)}$ is the same as if f is truly random. As a result, $|x_1 \cdots x_t\rangle\langle y_1 \cdots y_t|$ still has non-zero coefficient out only when each element of S appears even number of times in $(x_1,\ldots,x_t,y_1,\ldots,y_t)$. The rest of the proof is exactly the same as in Appendix A. Specifically for $t = 2$, since $\{|\psi_{S,f}\rangle\}$ is an $O(2^{-d})$-approximate 2-design, the indistinguishability from Haar random states follows from Theorem 3.4 and Theorem 3.11.

On the other hand, note that the states are given by

$$\frac{1}{2^{d/2}} \sum_{x \in S} (-1)^{f(x)} |x\rangle, \tag{4}$$

where d is the dimension of the subspace. For any cut of the qubits, (4) gives a tensor product decomposition of the state with at most 2^d terms. Hence, by Lemma 2.11, the Schmidt rank of the state is at most 2^d and the corresponding von Neumann entanglement entropy is at most d. □

Preparing phased subspace states. Subspace states are known to be efficiently preparable with $\mathcal{O}(nd)$ gates [KP22]. The 4-wise independent function $f : \{0,1\}^n \to \{0,1\}$ can be constructed from seeds of length $O(n)$ in $\mathrm{poly}(n)$ time (see e.g. [Vad12, Section 3.5]). The phases are put into the state using an efficient controlled operation.

4 Parallel-query secure PRU against local QNC^0

In previous sections we examined the pseudorandom properties of state designs against shallow quantum circuits. In this section we turn to unitary designs and show that they are also pseudorandom when queried in parallel, but against the more restricted class of circuits that are *geometrically local*. Specifically, we consider the distinguisher to be a circuit $C' \cdot (U^{\otimes t} \otimes \mathbb{I}) \cdot C$, where C, C' are both 1-dimensional geometrically local QNC^0 circuits, and U is either a unitary 2-design or a Haar random unitary applied on n consecutive qubits. The following folklore fact about geometrically local QNC^0 circuits is crucial for us:

Lemma 4.1. Let C be an 1-dimensional local QNC circuit of depth d on n qubits. Then for every $k \in [n-1]$, the Schmidt rank of the state $C|0^n\rangle$ between the first k qubits and the remaining $(n-k)$ qubits is at most 4^d.

Proof. We prove this by an induction over d, and the base case when $d = 0$ is trivial. Suppose that after the first d layer of gates we have the Schmidt decomposition $\sum_{i=1}^{4^d} \alpha_i |v_i\rangle |w_i\rangle$, where $|v_i\rangle$ is on k qubits and $|w_i\rangle$ is on $(n-k)$ qubits. In layer $d+1$, only the gate U (if it exists) that acts on the k-th and the $(k+1)$-th qubits would affect the Schmidt rank. We can write U with an arbitrary tensor product decomposition $U = \sum_{j=1}^{4} A_j \otimes B_j$ where $A_j, B_j \in \mathbb{C}^{2 \times 2}$, so that the state after applying U becomes

$$\sum_{i=1}^{4^d} \sum_{j=1}^{4} \alpha_i (A_j |v_i\rangle) \otimes (B_j |w_i\rangle).$$

By Lemma 2.11 we know that the above state has Schmidt rank at most 4^{d+1}, and it is not affected by the remaining gates in the same layer. $\square$

We will make use of a stronger statement that allows us to perform the Schmidt decomposition recursively on the t blocks of n qubits (a notion that we borrow from [GK12]):

Lemma 4.2 (Recursive Schmidt Decomposition). Let C be a 1-dimensional local QNC circuit of depth d on tn qubits. Then we can write the state $C|0^{tn}\rangle$ as

$$C|0^{tn}\rangle = \sum_{i_1,\dots,i_t=1}^{r} \alpha_{i_1,\dots,i_t} \cdot |v_{1,i_1}\rangle \otimes |v_{2,i_1,i_2}\rangle \otimes \cdots \otimes |v_{t,i_1,\dots,i_t}\rangle \tag{5}$$

where the Schmidt rank $r \le 4^d$, and $|v_{\tau,i_1,\dots,i_\tau}\rangle$ is an n-qubit state. For every $\tau \in [t]$, $i_1,\dots,i_\tau \in [r]$ and $i'_\tau \ne i_\tau$, we have the orthogonality condition

$$\langle v_{\tau,i_1,\dots,i_{\tau-1},i_\tau} | v_{\tau,i_1,\dots,i_{\tau-1},i'_\tau} \rangle = 0,$$

and the complex numbers $\alpha_{i_1,\dots,i_t}$ satisfy $\sum |\alpha_{i_1,\dots,i_t}|^2 = 1$.

Proof. Let r be the maximum Schmidt rank between the first τn and the remaining $(t-\tau)n$ qubits, for any $\tau \in [t]$. By Lemma 4.1 we have $r \le 4^d$. The recursive application of the Schmidt decomposition starts with the cut between the first n qubits and the remaining $(t-1)n$ qubits:

$$C|0^{tn}\rangle = \sum_{i_1=1}^{r} \alpha_{i_1} \cdot |v_{1,i_1}\rangle \otimes |w_{1,i_1}\rangle.$$

15

The next step is to perform a Schmidt decomposition over each $|w_{1,i_1}\rangle$, for which we need an upper bound on the Schmidt rank. We can also write $|w_{1,i_1}\rangle$ in the computational basis to get

$$C|0^{tn}\rangle = \sum_{i_1=1}^{r} \sum_{x\in\{0,1\}^n} \sum_{y\in\{0,1\}^{(t-2)n}} \alpha_{i_1}\beta_{i_1,x,y} \cdot |v_{1,i_1}\rangle|x\rangle|y\rangle.$$

Since $\{|v_{1,i_1}\rangle\}$ can be expanded into an orthonormal basis on the first n qubits, so can $\{|v_{1,i_1}\rangle|x\rangle\}$ on the first $2n$ qubits. As a result, the Schmidt rank of $C|0^{tn}\rangle$ between the first $2n$ and the remaining $(t-2)n$ qubits is exactly the rank of the $2^n r \times 2^{(t-2)n}$ matrix M, where

$$M((i_1,x),y) = \alpha_{i_1}\beta_{i_1,x,y}.$$

On the other hand, whenever $\alpha_{i_1} \neq 0$, the Schmidt rank of $|w_{1,i_1}\rangle$ on the same cut is the rank of the submatrix of $\alpha_{i_1}^{-1}M$, with the row index i_1 fixed, and thus is at most r. Therefore we get

$$C|0^{tn}\rangle = \sum_{i_1,i_2=1}^{r} \alpha_{i_1,i_2} \cdot |v_{1,i_1}\rangle \otimes |v_{2,i_1,i_2}\rangle \otimes |w_{2,i_1,i_2}\rangle.$$

Continuing the process for every $\tau \in [t]$ results in a tree-like structure as in (5), and the sum of squares of the coefficients is guaranteed to be 1 by the orthogonality of the states. $\square$

We also need the following lemma in analogy to Lemma 3.1, but on cross (off-diagonal) terms conjugated with Haar random unitaries.

Lemma 4.3. Let U be an n-qubit Haar random unitary, and let A be a subsystem with $|A| = k$ qubits and $B = [n] \setminus A$. For every two n-qubit states $|v\rangle$ and $|w\rangle$ such that $\langle v|w\rangle = 0$, we have

$$\mathbb{E}\left[\left\|\mathrm{Tr}_B[U|v\rangle\langle w|U^\dagger]\right\|_2^2\right] \leq 2^{k-n}.$$

Proof. Without loss of generality, we can assume that $U|w\rangle = |0^n\rangle$, and $U|v\rangle = |u\rangle$ is a uniformly random state orthogonal to $|0^n\rangle$. In this case, we can write

$$\mathrm{Tr}_B[U|v\rangle\langle w|U^\dagger] = \sum_{x\in\{0,1\}^{n-k}} (\mathbb{I}_A \otimes |x\rangle\langle x|_B)|u\rangle\langle 0^n|(\mathbb{I}_A \otimes |x\rangle\langle x|_B)$$

$$= (\mathbb{I}_A \otimes |0^{n-k}\rangle\langle 0^{n-k}|_B)|u\rangle\langle 0^n|.$$

And thus,

$$\left\|\mathrm{Tr}_B[U|v\rangle\langle w|U^\dagger]\right\|_2^2 = \mathrm{Tr}\left[(\mathbb{I}_A \otimes |0^{n-k}\rangle\langle 0^{n-k}|_B)|u\rangle\langle 0^n|0^n\rangle\langle u|(\mathbb{I}_A \otimes |0^{n-k}\rangle\langle 0^{n-k}|_B)\right]$$

$$= \mathrm{Tr}\left[(\mathbb{I}_A \otimes |0^{n-k}\rangle\langle 0^{n-k}|_B)|u\rangle\langle u|\right]$$

$$= \sum_{x\in\{0,1\}^k} \left|\langle x_A 0_B^{n-k}|u\rangle\right|^2.$$

As a side note, this directly implies that

$$\left\|\mathrm{Tr}_B[U|v\rangle\langle w|U^\dagger]\right\|_2 \leq 1 \tag{6}$$

regardless of the choice of $|v\rangle$, $|w\rangle$ or U, which will be useful later on.

Since $|u\rangle$ is a uniformly random state orthogonal to $|0^n\rangle$, for every $x \in \{0,1\}^n \setminus \{0^n\}$, $\langle x|u\rangle$ has the same distribution. Therefore every term in the summation above, except $|\langle 0^n|u\rangle|^2$, has the same expectation which is $1/(2^n - 1)$. Therefore we conclude that

$$\mathbb{E}\left[\left\|\text{Tr}_B[U|v\rangle\langle w|U^\dagger]\right\|_2^2\right] = \frac{2^k - 1}{2^n - 1} < 2^{k-n}.$$ □

The above result allows us to show that the output of non-adaptive queries of a Haar random unitary has the property that every subsystem of k qubits are almost maximally mixed, similar to Corollary 3.2, when the input state admits the recursive Schmidt decomposition.

Corollary 4.4. Let U be an n-qubit Haar random unitary. For $t \geq 1$, let $|\psi\rangle$ be a tn-qubit state that admits the recursive Schmidt decomposition with rank r as in Lemma 4.2. Let ρ_A be the reduced density matrix of $U^{\otimes t}|\psi\rangle\langle\psi|U^{\dagger\otimes t}$ over a subsystem A with $|A| = k$ qubits. Then for every $\delta > 0$, with probability at least $1 - O(r) \cdot 2^{k-n/2} \cdot \delta^{-1}$ over U, it holds that

$$\left\|\rho_A - \frac{1}{2^k}\mathbb{I}_k\right\|_1 \leq \delta.$$

Proof. To take the partial trace for subsystem A, we assume that A consists of $k_1, \ldots, k_t$ qubits in each block of n qubits respectively, and let B_τ be the part outside A in the τ-th block. With the recursive Schmidt decomposition (5) we can write

$$U^{\otimes t}|\psi\rangle\langle\psi|U^{\dagger\otimes t} = \sum_{\substack{i_1,\ldots,i_t, \\ j_1,\ldots,j_t=1}}^r \alpha_{i_1,\ldots,i_t}\overline{\alpha_{j_1,\ldots,j_t}} \cdot U|v_{1,i_1}\rangle\langle v_{1,j_1}|U^\dagger \otimes \cdots \otimes U|v_{t,i_1,\ldots,i_t}\rangle\langle v_{t,j_1,\ldots,j_t}|U^\dagger. \quad (7)$$

In addition, we can also write the decomposition (5) only to a certain level $\tau < t$ to get

$$|\psi\rangle = \sum_{i_1,\ldots,i_\tau=1}^r \alpha_{i_1,\ldots,i_\tau} \cdot |v_{1,i_1}\rangle \otimes |v_{2,i_1,i_2}\rangle \otimes \cdots \otimes |v_{\tau,i_1,\ldots,i_\tau}\rangle \otimes |w_{\tau,i_1,\ldots,i_\tau}\rangle,$$

where $|w_{\tau,i_1,\ldots,i_\tau}\rangle$ is a state on $(t-\tau)n$ qubits such that

$$\alpha_{i_1,\ldots,i_\tau}|w_{\tau,i_1,\ldots,i_\tau}\rangle = \sum_{i_{\tau+1},\ldots,i_t=1}^r \alpha_{i_1,\ldots,i_t} \cdot |v_{\tau+1,i_1,\ldots,i_{\tau+1}}\rangle \otimes \cdots \otimes |v_{t,i_1,\ldots,i_t}\rangle.$$

Notice that it also implies

$$|\alpha_{i_1,\ldots,i_\tau}|^2 = \sum_{i_{\tau+1},\ldots,i_t=1}^r |\alpha_{i_1,\ldots,i_t}|^2.$$

This way, we can group the summands in (7) depending on the smallest coordinate τ such that $i_\tau \neq j_\tau$ (when $\tau = t+1$, it means that $(i_1,\ldots,i_t)$ is identical to $(j_1,\ldots,j_t)$), and have

$$U^{\otimes t}|\psi\rangle\langle\psi|U^{\dagger\otimes t} = \sum_{\tau=1}^{t+1} \sum_{i_1,\ldots,i_\tau=1}^r \sum_{j_\tau \neq i_\tau} \alpha_{i_1,\ldots,i_\tau}\overline{\alpha_{i_1,\ldots,i_{\tau-1},j_\tau}} \cdot U|v_{1,i_1}\rangle\langle v_{1,i_1}|U^\dagger \otimes \cdots$$
$$\otimes U|v_{\tau,i_1,\ldots,i_\tau}\rangle\langle v_{\tau,i_1,\ldots,i_{\tau-1},j_\tau}|U^\dagger \otimes U^{\otimes(t-\tau)}|w_{\tau,i_1,\ldots,i_\tau}\rangle\langle w_{\tau,i_1,\ldots,i_{\tau-1},j_\tau}|U^{\dagger\otimes(t-\tau)}. \quad (8)$$

Now consider each summand in (8) with $\tau \leq t$. By Lemma 4.3, we have

$$\mathbb{E}\left[\left\|\text{Tr}_{B_\tau}[U|v_{\tau,i_1,\ldots,i_\tau}\rangle\langle v_{\tau,i_1,\ldots,i_{\tau-1},j_\tau}|U^\dagger]\right\|_2^2\right] \leq 2^{k_\tau - n}.$$

17

Therefore, after taking the partial trace, with the bound (6) on the Frobenius norms of the other blocks, we can bound the expected Frobenius norm of the entire summand by $\left|\alpha_{i_1,\ldots,i_\tau}\overline{\alpha_{i_1,\ldots,i_{\tau-1},j_\tau}}\right| \cdot 2^{(k_\tau-n)/2}$. Thus by linearity of expectation and triangular inequality, these summands in total has expected Frobenius norm of at most

$$\sum_{\tau=1}^{t}\sum_{i_1,\ldots,i_\tau=1}^{r}\sum_{j_\tau\neq i_\tau}\left|\alpha_{i_1,\ldots,i_\tau}\overline{\alpha_{i_1,\ldots,i_{\tau-1},j_\tau}}\right| \cdot 2^{(k_\tau-n)/2}$$

$$= \sum_{\tau=1}^{t}\sum_{i_1,\ldots,i_{\tau-1}=1}^{r}\left(\sum_{i_\tau=1}^{r}\left|\alpha_{i_1,\ldots,i_\tau}\right|\right)^2 \cdot 2^{(k_\tau-n)/2}$$

$$\leq \sum_{\tau=1}^{t}\sum_{i_1,\ldots,i_\tau=1}^{r}\left|\alpha_{i_1,\ldots,i_\tau}\right|^2 \cdot r \cdot 2^{(k_\tau-n)/2}$$

$$\leq r \cdot 2^{(k-n)/2}.$$

The remaining summands are those with $(i_1,\ldots,i_t) = (j_1,\ldots,j_t)$, and their sum is exactly

$$\sum_{i_1,\ldots,i_t=1}^{r}\left|\alpha_{i_1,\ldots,i_t}\right|^2 \cdot U|v_{1,i_1}\rangle\langle v_{1,i_1}|U^\dagger \otimes \cdots \otimes U|v_{t,i_1,\ldots,i_t}\rangle\langle v_{t,i_1,\ldots,i_t}|U^\dagger.$$

By Lemma 3.1, similarly to the deduction in Corollary 3.2, we know that the partial trace in each block, denoted by $M_\tau = \mathrm{Tr}_{B_\tau}[U|v_{\tau,i_1,\ldots,i_\tau}\rangle\langle v_{\tau,i_1,\ldots,i_\tau}|U^\dagger]$, satisfies that

$$\mathbb{E}\left[\left\|M_\tau - \frac{1}{2^{k_\tau}}\mathbb{I}_{k_\tau}\right\|_2\right] \leq 2^{(k_\tau-n)/2}.$$

Thus by a hybrid argument, we have

$$\mathbb{E}\left[\left\|M_1\otimes\cdots\otimes M_t - \frac{1}{2^k}\mathbb{I}_k\right\|_2\right] \leq \sum_{\tau=1}^{t}\mathbb{E}\left[\left\|M_1\otimes\cdots\otimes M_{\tau-1}\otimes\left(M_\tau - \frac{1}{2^{k_\tau}}\mathbb{I}_{k_\tau}\right)\right\|_2\right]$$

$$\leq \sum_{\tau=1}^{t}2^{(k_\tau-n)/2} \leq 2^{(k-n)/2}.$$

Since $\sum\left|\alpha_{i_1,\ldots,i_t}\right|^2 = 1$, we conclude that

$$\mathbb{E}\left[\left\|\rho_A - \frac{1}{2^k}\mathbb{I}_k\right\|_1\right] \leq 2^{k/2} \cdot \mathbb{E}\left[\left\|\rho_A - \frac{1}{2^k}\mathbb{I}_k\right\|_2\right] \leq (r+1) \cdot 2^{k-n/2}.$$

By Markov's inequality, we know that $\left\|\rho_A - \frac{1}{2^k}\mathbb{I}_k\right\|_1 \leq \delta$ holds with probability at least $1 - (r+1) \cdot 2^{k-n/2} \cdot \delta^{-1}$. □

Since the proof of Corollary 4.4 only uses the second moment properties of U (Lemma 3.1 and Lemma 4.3 to be exact), using Lemma 2.9 we can also conclude the following for approximate unitary 2-designs.

Corollary 4.5. Let $\{U_i\}$ be an ε-approximate unitary 2-design on n qubits. For $t \geq 1$, let $|\psi\rangle$ be a tn-qubit state that admits the recursive Schmidt decomposition with rank r as in Lemma 4.2. Let ρ_A be the reduced density matrix of $U_i^{\otimes t}|\psi\rangle\langle\psi|U_i^{\dagger\otimes t}$ over a subsystem A with $|A| = k$ qubits. Then for every $\delta > 0$, with probability at least $1 - 2^{O(k)} \cdot r \cdot (\varepsilon + 2^{-n})^{1/2} \cdot \delta^{-1}$ over i, it holds that

$$\left\|\rho_A - \frac{1}{2^k}\mathbb{I}_k\right\|_1 \leq \delta.$$

Combining the above corollaries together, we obtain the desired pseudorandomness against geometrically local QNC^0 circuits.

Theorem 4.6. Let $\{U_i\}$ be an ε-approximate unitary 2-design on n qubits for $\varepsilon = 2^{-\Omega(n)}$. Then $\{U_i\}$ is non-adaptive secure PRU against 1-dimensional geometrically local QNC circuits of depth $d = \log n - \omega(1)$. Moreover, the pre-processing part of the QNC circuit (before applying U_i) could have depth up to $o(n)$.

Proof. The proof follows from that of Theorem 3.4. Since the output of the QNC circuit depends on only $k = 2^d = o(n)$ qubits from the output of $U^{\otimes t}$, by Corollary 4.4 and Corollary 4.5 we know that the outputs have $\mathrm{negl}(n)$ trace distance between the Haar random U and unitary design $\{U_i\}$, with probability $1 - \mathrm{negl}(n)$ over the choice of the unitaries. This is because

$$2^{O(k)} \cdot r \cdot (\varepsilon + 2^{-n})^{1/2} \leq \mathrm{negl}(n)$$

when $k = o(n)$ and $\varepsilon = 2^{-\Omega(n)}$, and $r \leq 2^d$ from Lemma 4.2. In addition, the circuit depth d used to bound the Schmidt rank r only concerns the depth of the pre-processing part of the circuit, which can be raised up to $o(n)$. $\qquad\square$

Although Corollary 4.4 and Corollary 4.5 have virtually the same form as Corollary 3.2 and Corollary 3.3, we cannot use the techniques in Section 3.2 to directly obtain the similar security of PRU against QNC^0 circuits with AC^0 post-processing. The reason is that in Section 3.2 we have to limit the number of ancillae, and here they correspond to the qubits that U is not applied to, which we cannot put any natural restrictions on. However, in the artificial scenario where we force the the unitary to be applied in parallel on all qubits (and therefore allow no ancillae for the post-processing AC$^0 \circ$ QNC^0 circuit), we can obtain the following statement in analogy to Theorem 3.9:

Theorem 4.7. Let $\{U_i\}$ be an ε-approximate unitary 2-design on n qubits for $\varepsilon = 2^{-\log^{\omega(1)} n}$. Then $\{U_i\}$ is PRU against a subclass of AC$^0 \circ$ QNC^0 circuits on a multiple of n qubits, where U_i is applied non-adaptively over all the qubits used by the circuit, and the pre-processing QNC^0 circuit before applying U_i is 1-dimensional geometrically local.

5 Discussion and outlook

Our work initiates the study of unconditionally fooling shallow quantum circuits. In general, statistical and computational notions of quantum pseudorandomness—such as t-designs and PRS—are incomparable. Our work shows that in the low-complexity regime, the two notions can in fact overlap and illustrate rich connections to computational complexity theory, reminiscent of the intimate relation between hardness and randomness in classical computation. Our work leaves a number of interesting questions regarding the connections between hardness and quantum pseudorandomness. We discuss a few of these below.

Fooling stronger circuits. What are the strongest class of quantum circuits that t-designs, or in particular 2-designs, can fool? We conjecture that 2-designs are computationally secure against QAC^0 circuits as well. To prove this, it suffices to show a QAC analogy of Braverman's result on *almost k-wise maximally mixed* input states. These states have the property that every subsystem on k qubits is close to being maximally mixed, and we conjecture that they cannot be distinguished from the true maximally mixed state by QAC circuits of small depths.

A more immediate improvement on our results would be the removal of the constraints on ancillae in Theorem 3.11 and Corollary 3.12. The reason we require a bounded number of ancillae is purely technical: we need this to argue that the output distribution has high min-entropy, as otherwise the k-wise indistinguishability would not guarantee that we can fool AC^0 circuits. We note that a similar techinical issue occured in the $AC^0 \circ QNC^0$ lower bound lower bound result of [Slo24].

Stronger security for PRU. The unconditional security we proved for PRU is quite limited. Specifically, in Theorem 4.6 we could only show security when the unitaries are queried non-adaptively, while the adversaries are 1-dimensional geometrically local QNC circuits. Can we lift the requirement of non-adaptivity or geometric locality?

We conjecture that new constructions are necessary in order to achieve security against adaptive queries. In other words, there exist (approximate) t-designs that are not PRU against QNC^0 circuits with adaptive queries. Such an example that works for arbitrary $t \leq \text{poly}(n)$ would also give a separation between non-adaptive and adaptive PRU.

Optimal pseudoentanglement. It is known that the optimal entanglement entropy gap for pseudoentanglement is $\omega(\log n)$ versus $O(n)$, which is achievable across every cut when assuming the existence of post-quantum one way functions [ABF+24]. In Corollary 3.13 we showed explicit examples of unconditional pseudoentanglement, which is optimal across every cut against QNC^0 distinguishers, but only $\log^{\omega(1)} n$ versus $O(n)$ against an $AC^0 \circ QNC^0$ adversary. Can we also construct unconditional optimal pseudoentanglement against $AC^0 \circ QNC^0$, or even stronger circuits?

PRG against shallow circuits. In this work we did not consider classical pseudorandom primitives, such as PRGs. One of the reasons is that we need to be careful about the definition when the adversaries are shallow quantum circuits, since it makes a difference whether or not we allow access to multiple copies of the PRG output. In fact, if the adversary could only access a single copy, then the classical Nisan-Wigderson generator [Nis92; NW94] instantiated by the parity function would directly give an unconditional PRG with $\text{polylog}(n)$ seed length against QAC^0 circuits with bounded number of ancillae, using the recently results on the hardness of parity against QAC^0 [NPV+24; ADO+24]. However, the hardness proofs in these works, which examine the Pauli spectrum of the QAC^0 circuit, break down when allowing multi-copy access to the input, as even a single classical fan-out gate could significantly increase the Pauli weight of the overall circuit.

As a result, designing unconditional multi-copy secure PRG against QAC^0 circuits remains to be an intriguing open problem. A follow-up direction is to use such PRG to construct other pseudorandom primitives such as pseudorandom functions or even PRS and PRU. Notice that we have recipes for these constructions against polynomial-sized adversaries, but in order to work against bounded adversaries, the construction needs to be super-efficient and such constructions are largely unknown.

Fooling other models. Finally, an open-ended question is to find unconditional pseudorandom objects against other restricted models of quantum computation, with or without our constructions and techniques. It was shown in [GR22] that the INW generator [INW94] is secure against space-bounded quantum computation. To study PRS or PRU against such models, the first challenge lies in finding the most relevant and useful definition, which we leave for future work.

Acknowledgements

References

[ABF⁺24]   Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. "Quantum Pseudoentanglement". Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 2:1–2:21, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.

[ABC⁺24]   Chris Akers, Adam Bouland, Lijie Chen, Tamara Kohler, Tony Metger, and Umesh Vazirani. *Holographic pseudoentanglement and the complexity of the ads/cft dictionary*, 2024. arXiv: 2411.04978 [hep-th].

[AGQ⁺23]   Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. *Pseudorandom (function-like) quantum state generators: new definitions and applications*, 2023. arXiv: 2211.01444 [quant-ph].

[AQY22]   Prabhanjan Ananth, Luowen Qian, and Henry Yuen. *Cryptography from pseudorandom quantum states*, 2022. arXiv: 2112.10020 [quant-ph].

[ADO⁺24]   Anurag Anshu, Yangjing Dong, Fengning Ou, and Penghui Yao. *On the computational power of qac0 with barely superlinear ancillae*, 2024. arXiv: 2410.06499 [quant-ph].

[AD25]   Srinivasan Arunachalam and Arkopal Dutt. "Polynomial-time tolerant testing stabilizer states". *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, pages 1234–1241. ACM, June 2025.

[BIV⁺16]   Andrej Bogdanov, Yuval Ishai, Emanuele Viola, and Christopher Williamson. "Bounded indistinguishability and the complexity of recovering secrets". *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference*, volume 9816 of *Lecture Notes in Computer Science*, pages 593–618. Springer, 2016.

[BV10]   Andrej Bogdanov and Emanuele Viola. "Pseudorandom bits for polynomials". *SIAM J. Comput.*, 39(6):2464–2486, 2010.

[BFG⁺23]   Adam Bouland, Bill Fefferman, Soumik Ghosh, Tony Metger, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. *Public-key pseudoentanglement and the hardness of learning ground state entanglement structure*, 2023. arXiv: 2311.12017 [quant-ph].

[BZZ24]   Adam Bouland, Chenyi Zhang, and Zixin Zhou. *On the hardness of learning ground state entanglement of geometrically local hamiltonians*, 2024. arXiv: 2411.04353 [quant-ph].

[BHH16a]   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.

[BHH16b]   Fernando G.S.L. Brandão, Aram W. Harrow, and Michał Horodecki. "Efficient quantum pseudorandomness". *Physical Review Letters*, 116(17), April 2016.

[Bra08]   Mark Braverman. "Polylogarithmic independence fools AC^0 circuits". *J. ACM*, 57(5), June 2008.

[CCG⁺25] Shantanav Chakraborty, Soonwon Choi, Soumik Ghosh, and Tudor Giurgică-Tiron. *Fast computational deep thermalization*, 2025. arXiv: 2507.13670 [quant-ph].

[CSB⁺25] Laura Cui, Thomas Schuster, Fernando Brandao, and Hsin-Yuan Huang. *Unitary designs in nearly optimal depth*, 2025. arXiv: 2507.06216 [quant-ph].

[DHB22] Alexander M Dalzell, Nicholas Hunter-Jones, and Fernando GSL Brandão. "Random quantum circuits anticoncentrate in log depth". *PRX Quantum*, 3(1):010333, 2022.

[DET⁺10] Anindya De, Omid Etesami, Luca Trevisan, and Madhur Tulsiani. "Improved pseudo-random generators for depth 2 circuits". *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2010*, volume 6302 of *Lecture Notes in Computer Science*, pages 504–517. Springer, 2010.

[FI24] Xiaozhou Feng and Matteo Ippoliti. *Dynamics of pseudoentanglement*, 2024. arXiv: 2403.09619 [quant-ph].

[FI25] Xiaozhou Feng and Matteo Ippoliti. "Dynamics of pseudoentanglement". *Journal of High Energy Physics*, 2025(2), February 2025.

[GK12] Sevag Gharibian and Julia Kempe. "Approximation algorithms for QMA-complete problems". *SIAM J. Comput.*, 41(4):1028–1050, 2012.

[GR22] Uma Girish and Ran Raz. "Eliminating intermediate measurements using pseudorandom generators". *13th Innovations in Theoretical Computer Science Conference, ITCS 2022*, volume 215 of *LIPIcs*, 76:1–76:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[GB23] Tudor Giurgica-Tiron and Adam Bouland. *Pseudorandomness from subset states*, 2023. arXiv: 2312.09206 [quant-ph].

[GMR⁺12] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil P. Vadhan. "Better pseudorandom generators from milder pseudorandom restrictions". *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012*, pages 120–129. IEEE Computer Society, 2012.

[GE24] Manuel Goulão and David Elkouss. *Pseudo-entanglement is necessary for efi pairs*, 2024. arXiv: 2406.06881 [quant-ph].

[GNW21] David Gross, Sepehr Nezami, and Michael Walter. "Schur–weyl duality for the clifford group with applications: property testing, a robust hudson theorem, and de finetti representations". *Communications in Mathematical Physics*, 385(3):1325–1393, June 2021.

[GLG⁺24] Andi Gu, Lorenzo Leone, Soumik Ghosh, Jens Eisert, Susanne F. Yelin, and Yihui Quek. "Pseudomagic quantum states". *Physical Review Letters*, 132(21), May 2024.

[HS19] Prahladh Harsha and Srikanth Srinivasan. "On polynomial approximations to AC^0". *Random Struct. Algorithms*, 54(2):289–303, 2019.

[HH24] Pooya Hatami and William Hoza. "Paradigms for unconditional pseudorandom generators". *Found. Trends Theor. Comput. Sci.*, 16(1-2):1–210, 2024.

[HVV06] Alexander Healy, Salil P. Vadhan, and Emanuele Viola. "Using nondeterminism to amplify hardness". *SIAM J. Comput.*, 35(4):903–931, 2006.

[Hoz21] William M. Hoza. "Better pseudodistributions and derandomization for space-bounded computation". *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021*, volume 207 of *LIPIcs*, 28:1–28:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[INW94]   Russell Impagliazzo, Noam Nisan, and Avi Wigderson. "Pseudorandomness for network algorithms". *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, pages 356–364. ACM, 1994.

[JMW24]   Fernando Granha Jeronimo, Nir Magrafta, and Pei Wu. *Pseudorandom and pseudoentangled states from subset states*, 2024. arXiv: 2312.15285 [quant-ph].

[JLS18]   Zhengfeng Ji, Yi-Kai Liu, and Fang Song. *Pseudorandom quantum states*. *Advances in Cryptology – CRYPTO 2018*. Springer International Publishing, 2018, pages 126–152.

[KP22]   Iordanis Kerenidis and Anupam Prakash. *Quantum machine learning with subspace states*, 2022. arXiv: 2202.00054 [quant-ph].

[LLZ+18]   Zi-Wen Liu, Seth Lloyd, Elton Zhu, and Huangjun Zhu. "Entanglement, quantum randomness, and complexity beyond scrambling". *Journal of High Energy Physics*, 2018(7):1–62, 2018.

[Lub78]   Elihu Lubkin. "Entropy of an n-system from its correlation with a k-reservoir". *Journal of Mathematical Physics*, 19(5):1028–1031, 1978.

[Lub94]   Michael George Luby. *Pseudorandomness and Cryptographic Applications*. Princeton University Press, USA, 1994.

[Mel24]   Antonio Anna Mele. "Introduction to haar measure tools in quantum information: a beginner's tutorial". *Quantum*, 8:1340, 2024.

[MPS+24]   Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. *Pseudorandom unitaries with non-adaptive security*, 2024. arXiv: 2402.14803 [quant-ph].

[NPV+24]   Shivam Nadimpalli, Natalie Parham, Francisca Vasconcelos, and Henry Yuen. "On the Pauli spectrum of QAC0". *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, pages 1498–1506, Vancouver, BC, Canada. Association for Computing Machinery, 2024.

[NN93]   Joseph Naor and Moni Naor. "Small-bias probability spaces: efficient constructions and applications". *SIAM J. Comput.*, 22(4):838–856, 1993.

[NC10]   Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010.

[Nis91]   Noam Nisan. "Pseudorandom bits for constant depth circuits". *Combinatorica*, 11(1):63–70, March 1991.

[Nis92]   Noam Nisan. "Pseudorandom generators for space-bounded computation". *Combinatorica*, 12(4):449–461, 1992.

[NW94]   Noam Nisan and Avi Wigderson. "Hardness vs randomness". *J. Comput. Syst. Sci.*, 49(2):149–167, 1994.

[NZ96]   Noam Nisan and David Zuckerman. "Randomness is linear in space". *J. Comput. Syst. Sci.*, 52(1):43–52, 1996.

[Reg24]   Oded Regev. *On lattices, learning with errors, random linear codes, and cryptography*, 2024. arXiv: 2401.03703 [cs.CR].

[SZ99]   Michael E. Saks and Shiyu Zhou. "$\mathsf{BP_HSPACE}(S) \subseteq \mathsf{DSPACE}(S^{3/2})$". *J. Comput. Syst. Sci.*, 58(2):376–403, 1999.

[SHH25]   Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. "Random unitaries in extremely low depth". *Science*, 389(6755):92–96, July 2025.

[Slo24]   Joseph Slote. "Parity vs. AC0 with simple quantum preprocessing". *15th Innovations in Theoretical Computer Science Conference, ITCS 2024*, volume 287 of *LIPIcs*, 92:1–92:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[Tal17]   Avishay Tal. "Tight bounds on the fourier spectrum of AC^0". *32nd Computational Complexity Conference, CCC 2017*, volume 79 of *LIPIcs*, 15:1–15:31. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[Vad12]   Salil P Vadhan. "Pseudorandomness". *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012.

[Vio09]   Emanuele Viola. "The sum of D small-bias generators fools polynomials of degree D". *Comput. Complex.*, 18(2):209–217, 2009.

[YE23]   Lisa Yang and Netta Engelhardt. *The complexity of learning (pseudo)random dynamics of black holes and other chaotic systems*, 2023. arXiv: 2302.11013 [hep-th].

[Zha21]   Mark Zhandry. "How to construct quantum random functions". *J. ACM*, 68(5), August 2021.

A   Design properties of random phased subspace states

Here we show that the random phased subspace states $|\psi_{S,f}\rangle$, defined in Definition 3.6, form an approximate design. Fixing the subspace S and taking the average over the random function $f : S \to \{0, 1\}$, we have

$$\mathbb{E}_f \left[|\psi_{S,f}\rangle\langle\psi_{S,f}|^{\otimes t} \right] = \frac{1}{2^{dt}} \sum_{\substack{x_1,\dots,x_t \in S \\ y_1,\dots,y_t \in S}} \mathbb{E}_f \left[(-1)^{f(x_1)+\dots+f(x_t)+f(y_1)+\dots+f(y_t)} \right] |x_1 \cdots x_t\rangle\langle y_1 \cdots y_t|, \quad (9)$$

where the term $|x_1 \cdots x_t\rangle\langle y_1 \cdots y_t|$ has non-zero coefficient only when each element of S appears an even number of times in $(x_1, \dots, x_t, y_1, \dots, y_t)$. Among those let us consider the ones such that $x_1, \dots, x_t$ are all distinct (so that $y_1, \dots, y_t$ is a permutation of $x_1, \dots, x_t$); the partial summation over these terms is

$$\frac{1}{2^{dt}} \sum_{\substack{x_1,\dots,x_t \in S \\ x_i \neq x_j, \forall i < j \\ \pi \in \mathcal{S}_t}} |x_1 \cdots x_t\rangle\langle x_{\pi(1)} \cdots x_{\pi(t)}|$$

$$= \frac{1}{2^{dt}} \sum_{\{x_1,\dots,x_t\} \subset S} \left(\sum_{\pi \in \mathcal{S}_t} |x_{\pi(1)} \cdots x_{\pi(t)}\rangle \right) \left(\sum_{\pi \in \mathcal{S}_t} \langle x_{\pi(1)} \cdots x_{\pi(t)}| \right)$$

$$= \frac{t!}{2^{dt}} \sum_{X \subset S, |X|=t} |\mathsf{Sym}_X\rangle\langle\mathsf{Sym}_X|.$$

Here $\mathcal{S}_t$ is the symmetric group on $[t]$, and $|\mathsf{Sym}_X\rangle = \frac{1}{\sqrt{t!}} \sum_{\pi \in \mathcal{S}_t} |x_{\pi(1)} \cdots x_{\pi(t)}\rangle$ for $X = \{x_1, \dots, x_t\}$. As a result, the above partial sum has trace

$$\frac{t!}{2^{dt}} \cdot \binom{2^d}{t} = \frac{2^d \cdot (2^d - 1) \cdots \cdots (2^d - t + 1)}{2^{dt}} \geq 1 - \frac{t^2}{2^d}.$$

Also notice that the partial sum can be thought as the projection of $\mathbb{E}_f \left[|\psi_{S,f}\rangle\langle\psi_{S,f}|^{\otimes t} \right]$ onto the subspace spanned by $\{|\mathsf{Sym}_X\rangle\}_{X \subset S, |X|=t}$, implying that the remaining part is still positive-definite.

This allows us to bound the trace distance:

$$\left\| \mathbb{E}_f \left[|\psi_{S,f}\rangle\langle\psi_{S,f}|^{\otimes t} \right] - \binom{2^d}{t}^{-1} \sum_{X \subset S, |X|=t} |\mathsf{Sym}_X\rangle\langle\mathsf{Sym}_X| \right\|_1$$

$$\leq \left\| \mathbb{E}_f \left[|\psi_{S,f}\rangle\langle\psi_{S,f}|^{\otimes t} \right] - \frac{t!}{2^{dt}} \sum_{X \subset S, |X|=t} |\mathsf{Sym}_X\rangle\langle\mathsf{Sym}_X| \right\|_1 + \left| \frac{t!}{2^{dt}} \cdot \binom{2^d}{t} - 1 \right| \leq \frac{2t^2}{2^d}.$$

Now we think of S as a uniformly random d-dimensional subspace of $\{0,1\}^n$. For a uniformly random $X \subset S$ with $|X| = t$, the elements in X are linearly dependent with probability at most

$$\frac{1}{2^d} + \frac{1}{2^{d-1}} + \cdots + \frac{1}{2^{d-t+1}} < \frac{1}{2^{d-t}}.$$

Similarly, when X is a uniformly random size-t subset of $\{0,1\}^n$, the elements in X are linearly dependent with probability at most 2^{t-n}. When conditioned on linear independence, the distributions of X in both cases are the same, and thus

$$\left\| \mathbb{E}_S \left[\binom{2^d}{t}^{-1} \sum_{X \subset S, |X|=t} |\mathsf{Sym}_X\rangle\langle\mathsf{Sym}_X| \right] - \binom{2^n}{t}^{-1} \sum_{X \subset \{0,1\}^n, |X|=t} |\mathsf{Sym}_X\rangle\langle\mathsf{Sym}_X| \right\|_1 \leq 2^{t-d} + 2^{t-n}.$$

Hence we conclude that

$$\left\| \mathbb{E}_{S,f} \left[|\psi_{S,f}\rangle\langle\psi_{S,f}|^{\otimes t} \right] - \binom{2^n}{t}^{-1} \sum_{X \subset \{0,1\}^n, |X|=t} |\mathsf{Sym}_X\rangle\langle\mathsf{Sym}_X| \right\|_1 \leq \frac{2t^2}{2^d} + 2^{t-d} + 2^{t-n} = O(2^{t-d}).$$

On the other hand, for the Haar random state $|\psi_{\mathsf{Haar}}\rangle$, it is well known that (see e.g. [Mel24])

$$\mathbb{E}\left[|\psi_{\mathsf{Haar}}\rangle\langle\psi_{\mathsf{Haar}}|^{\otimes t} \right] = \binom{2^n + t - 1}{t}^{-1} \Pi_{\mathsf{sym}}^{(t,2^n)},$$

where $\Pi_{\mathsf{sym}}^{(t,2^n)}$ is the projection onto the symmetric subspace of $(\mathbb{C}^{2^n})^{\otimes t}$. Since all the $|\mathsf{Sym}_X\rangle$ take up $\binom{2^n}{t}$ dimensions in the subspace, their weight in $\Pi_{\mathsf{sym}}^{(t,2^n)}$ is at least

$$\binom{2^n}{t} \Big/ \binom{2^n + t - 1}{t} = \frac{2^n \cdot (2^n - 1) \cdot \cdots \cdot (2^n - t + 1)}{2^n \cdot (2^n + 1) \cdot \cdots \cdot (2^n + t - 1)} \geq 1 - \frac{t^2}{2^n}.$$

This means that

$$\left\| \mathbb{E}\left[|\psi_{\mathsf{Haar}}\rangle\langle\psi_{\mathsf{Haar}}|^{\otimes t} \right] - \binom{2^n}{t}^{-1} \sum_{X \subset \{0,1\}^n, |X|=t} |\mathsf{Sym}_X\rangle\langle\mathsf{Sym}_X| \right\|_1 \leq \frac{2t^2}{2^n},$$

and we can finally obtain that

$$\left\| \mathbb{E}_{S,f} \left[|\psi_{S,f}\rangle\langle\psi_{S,f}|^{\otimes t} \right] - \mathbb{E}\left[|\psi_{\mathsf{Haar}}\rangle\langle\psi_{\mathsf{Haar}}|^{\otimes t} \right] \right\|_1 \leq \frac{2t^2}{2^n} + O(2^{t-d}) \leq O(2^{t-d}).$$